arXiv:1612.09339v1 [cs.SE] 29 Dec 2016

A Rigorous Framework for Specification, Analysis and Enforcement of Access Control Policies

Andrea Margheri[1], Massimiliano Masi[2], Rosario Pugliese[3], and Francesco Tiezzi[4]

[1]Electronics and Computer Science, University of Southampton
[2]Tiani "Spirit" GmbH
[3]Università degli Studi di Firenze
[4]Università di Camerino

Abstract

Access control systems are widely used means for the protection of computing systems. They are defined in terms of access control policies regulating the accesses to system resources. In this paper, we introduce a formally-defined, fully-implemented framework for specification, analysis and enforcement of attribute-based access control policies. The framework rests on FACPL, a language with a compact, yet expressive, syntax for specification of real-world access control policies and with a rigorously defined denotational semantics. The framework enables the automatic verification of properties regarding both the authorisations enforced by single policies and the relationships among multiple policies. Effectiveness and performance of the analysis rely on a semantic-preserving representation of FACPL policies in terms of SMT formulae and on the use of efficient SMT solvers. Our analysis approach explicitly addresses some crucial aspects of policy evaluation, as e.g. missing attributes, erroneous values and obligations, which are instead overlooked in other proposals. The framework is supported by Java-based tools, among which an Eclipse-based IDE offering a tailored development and analysis environment for FACPL policies and a Java library for policy enforcement. We illustrate the framework and its formal ingredients by means of an e-Health case study, while its effectiveness is assessed by means of performance stress tests and experiments on a well-established benchmark.

1 Introduction

Nowadays computing systems have pervaded every daily activity and prompted the proliferation of several innovative services and applications. These modern distributed systems manage a huge amount of data that, due to its importance and societal impact, has brought out security issues of paramount importance. Controlling the access to system resources is thus crucial to prevent unauthorised accesses that could jeopardise trustworthiness of data.

This has prompted an increasing research interest towards access control systems, which are the first line of defence for the protection of computing systems. They are defined by *rules* that establish under which conditions a subject's *request* for accessing a resource has to be permitted or denied. In practice, it amounts to restrict physical and logical access rights of subjects to system resources.

Access control is a broad field, covering several different approaches, using different technologies and involving various degrees of complexity. Since the first applications in operating systems, to the more recent ones in distributed systems, many access control approaches have been proposed. Traditional approaches are based on the identity of subjects, either directly – e.g., Access Control Matrix [29] – or through predefined features, such as roles or groups – e.g., Role-Based Access Control (RBAC [17]). These approaches are however inadequate for dealing with modern distributed systems, as they suffer from scalability and interoperability issues. Moreover, they cannot easily encompass information representing the evaluation context, as e.g. system status or current time. An alternative approach that permits to overcome these problems is Attribute-Based Access Control (ABAC) [21]. Here, the rules are based on *attributes*, which represent arbitrary security-relevant information exposed by the system, the involved subjects, the action to be performed,

or by any other entity of the evaluation context relevant to the rules at hand. Thus, ABAC permits defining fine-grained, flexible and context-aware access control rules that are expressive enough to uniformly represent all the other approaches [25]. Attribute-based rules are typically hierarchically structured and paired with strategies for resolving possible conflicting authorisation results. These structured specifications are called *policies*; from this name derives the terminology Policy-Based Access Control (PBAC) [37], sometimes used in place of ABAC.

Many languages have been proposed for the specification of access control policies (see, e.g., [19] for a survey). Among the proposed languages, in the authors' knowledge, the OASIS standard eXtensible Access Control Markup Language (XACML) [38] is the best-known one. Due to its XML-based syntax and the advanced access control features it provides, XACML is commonly used in many real-world systems, e.g., in service-oriented ones. However, the management of real access control policies is in practice cumbersome and error-prone, and should be supported by rigorous analysis techniques. Unfortunately, XACML is generally acknowledged as lacking of a formally defined semantics (see, e.g., [41, 8, 40, 2]), which makes it difficult the specification and realisation of analysis techniques.

To tackle these difficulties, we introduce a formally-defined, fully-implemented framework, based on Formal Access Control Policy Language (FACPL), supporting developers in the specification, analysis and enforcement of access control policies.

The FACPL-based Access Control Framework

The FACPL language defines a core, yet expressive, syntax for specification of high-level access control policies. It is inspired by XACML (with which it shares the main traits of the policy structure and some terminology), but it refines some aspects of XACML and introduces novel features from the access control literature. Evaluation of FACPL policies is formalised by a denotational semantics, which clarifies intricate aspects of access controls like, e.g., management of missing attributes (i.e. attributes controlled by a policy but not provided by the request to authorise) and formalisation of combining algorithms (i.e. strategies to resolve conflictual decisions that policy evaluation can generate).

The analysis functionalities offered by our framework enable verification of two main groups of properties of FACPL policies. The *authorisation properties* permit to statically reason on the result of the evaluation of a policy with respect to a specific request, by also considering additional attributes that can be possibly introduced in the request at run-time and that might lead to unexpected authorisations. Instead, the *structural properties* permit to statically reason on the whole set of results of the evaluation of one or more policies and can be exploited, e.g., to implement maintenance and *change-impact analysis* [18] techniques.

The verification of these properties requires extensive checks on very large (possibly infinite) amounts of requests, hence support through software tools is essential. As no off-the-shelf analysis tool directly takes FACPL specifications in input, our framework exploits a constraint formalism that permits uniformly representing policy elements and enabling automated analysis. The constraint formalism we introduce is based on Satisfiability Modulo Theories (SMT) formulae, that is formulae defining satisfiability problems involving multiple theories, e.g. boolean and linear arithmetic ones. The relevant progress made in the development of automatic SMT solvers has led SMT to be extensively employed in diverse analysis applications [13], even for access control policies [2, 46]. In practice, SMT-based approaches are more effective than many other ones, like e.g. decision diagrams [18] or description logic [26]. The soundness of our analysis techniques is guaranteed by the correspondence, which we formally prove, between the semantics of FACPL policies and that of their constraint-based representations.

Our framework is supported by a Java-based software *toolchain*. The key software tool is an Eclipse-based IDE that offers a tailored development and analysis environment for FACPL policies. Specifically, it helps access control policy developers in the tasks of specification, analysis and enforcement of policies by providing, e.g., static checks on the code and automatic generation of runnable SMT and Java code. The evaluation of the SMT code relies on the Z3 solver [12], while the policy enforcement relies on an expressly developed Java library.

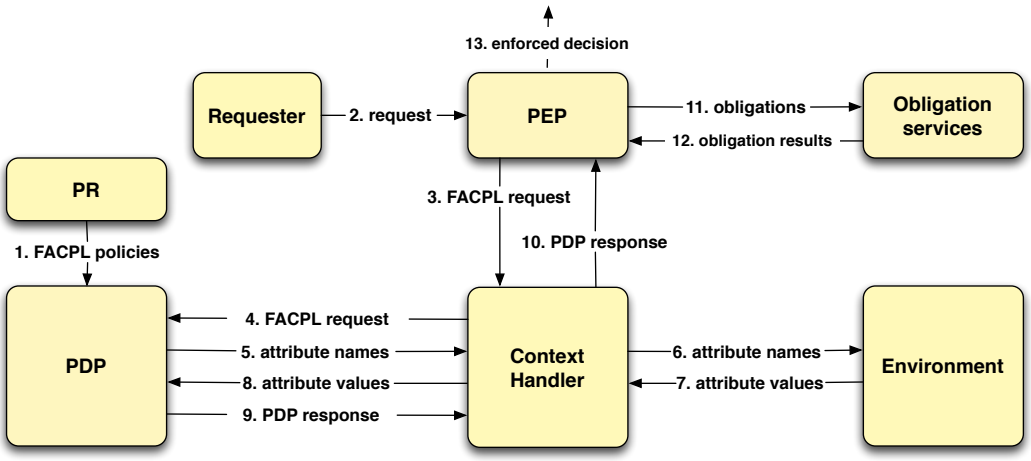

Fig. 1: The FACPL evaluation process

Contributions

The main contribution of this paper is the development of a comprehensive methodology supporting the whole life-cycle of access control policies, from their specification and analysis to their enforcement. Each ingredient of the methodology is formally introduced in this paper, together with its tool implementation. The tools allow access control system developers to use formally-defined functionalities without requiring them to be familiar with formal methods.

Our methodology enhances the proposals from the literature to different extents, in order to provide a single framework where all the relevant functionalities are expressed and formalised in a uniform manner. Indeed, XACML does not come with any formal specification and, hence, analysis; the formally-grounded proposals in [24, 8, 40, 10] do not offer any supporting tools; the SMT-based analysis proposals in [2, 46] do not support some crucial features (e.g., missing attributes and obligation instantiation). Detailed comparisons with the relevant literature are in Section 9.

Our aim is to design an expressive language whose formal foundations enable tool-supported analysis, rather than to face XACML semantic issues or supersede it. Further contributions of this paper are summarised below.

- The FACPL semantics manages missing attributes in a way similar to [8, 10] and extends it with explicit error management.
- The formalisation of combining algorithms extends that of [31] with explicit combination of obligations and with different instantiation strategies.
- The authorisation properties explicitly take into account the non-monotonicity issue of policy evaluation [45] by appropriately employing the request extensions set of [9] for property formalisation.
- The main structural properties of [18] and [26] are uniformly formalised in terms of policy semantics.
- The constraint formalism defines a low-level, tool-independent representation of attribute-based policies that is capable to deal with all issues regarding policy evaluation.
- The validation of the proposal is carried out through experiments on a standard benchmark in the field of access control tools, i.e. the CONTINUE [28] case study.

This paper is a revised and extended version of [36, 34]. Besides significant revisions and extensions of syntax and semantics of FACPL (we refer to Section 9 for a detailed comparison) this paper proposes a complete development methodology for access control policies. Most of all, differently from previous

works, we introduce a constraint-based representation of FACPL policies enabling the verification of various properties through SMT solvers.

Summary of the rest of the paper. In Section 2 we overview the FACPL evaluation process. In Section 3 we introduce an e-Health case study we use throughout the paper as a running example. In Section 4 we present the syntax of FACPL and its informal semantics, together with the FACPL-based specification of the case study. In Section 5 we formally define the FACPL semantics. In Section 6 we introduce the constraint formalism and the representation it enables of FACPL policies. In Section 7 we introduce various properties for access control policies and their verification via SMT solvers. In Section 8 we outline the Java-based software toolchain. In Section 9 we discuss the closest related work and, finally, in Section 10 we conclude and touch upon directions for future work. Appendixes A and B report, respectively, all the definitions for combining algorithms, and the proofs of the formal results.

2 The FACPL Evaluation Process

The FACPL evaluation process of (access control) policies and requests is shown in Figure 1. It defines the interactions, leading to the final authorisation decision, among three key components: the Policy Repository (PR), the Policy Decision Point (PDP) and the Policy Enforcement Point (PEP). These entities and their interactions were introduced in [49] to define the evaluation process of policy-based systems. Each policy language, e.g. XACML, has then tailored them according to its specific features.

The evaluation process assumes that system resources are paired with one or more FACPL policies, which define the credentials necessary to gain access to such resources. The PR stores the policies and supplies them to the PDP (step 1), which then decides if the access can be granted.

When PEP receives a request (step 2), the credentials contained in the request are encoded as a sequence of *attribute* elements (i.e., name-value pairs representing arbitrary information relevant for evaluating the access request) forming a FACPL request (step 3). PEPs can have many different forms, e.g. a gateway or a Web server. Therefore, this encoding allows policies and requests to be written and evaluated independently of their specific nature.

The *context handler* sends the request to the PDP (step 4), by possibly adding environmental attributes, e.g. request receiving time, that may be used in the evaluation.

The PDP *authorisation process* computes the *PDP response* for the request by checking the attributes, that may belong either to the request or to the environment (steps 5-8), against the controls contained in the policies. The PDP response (steps 9-10) contains an authorisation *decision* and possibly some *obligations*.

The decision is one among permit, deny, not-app and indet[1]. The meaning of the first two ones is obvious, the third one means that there is no policy that applies to the request and the latter one means that some errors have occurred during the evaluation. Policies can automatically manage these errors by using operators that combine, according to different strategies, indet decisions with the others.

Obligations are instead additional actions connected to the access control system that must be discharged by the PEP through appropriate *obligation services* (steps 11-12). Obligations usually correspond to, e.g., updating a log file, sending a message or executing a command. The *enforcement process* performed by the PEP determines the *enforced decision* (step 13) on the basis of the result of obligations discharge. This decision could differ from that of the PDP and is the outcome of the evaluation process.

It is worth noticing that the FACPL evaluation process guarantees separation of concerns among policies, their evaluation and the system itself. Among others, the main advantages it ensures are: (i) different types of requests can be handled, as the PEP can appropriately encode them in the format required by the PDP; (ii) the PDP can be placed in any point of the system, with the PEP acting as a gateway or a proxy; (iii) the PR can be also instantiated to support dynamic, possibly regulated, modifications of policies[2].

[1] The FACPL supporting tools can handle the same extended indeterminate values dealt with by XACML (see Section 8). However, for the sake of presentation, in the formal specification of FACPL we only consider a single indeterminate value, rather than the whole set.

[2] When PR provides also support for the specification of administration controls on policy modifications, it is usually called *Policy Administration Point* (PAP).

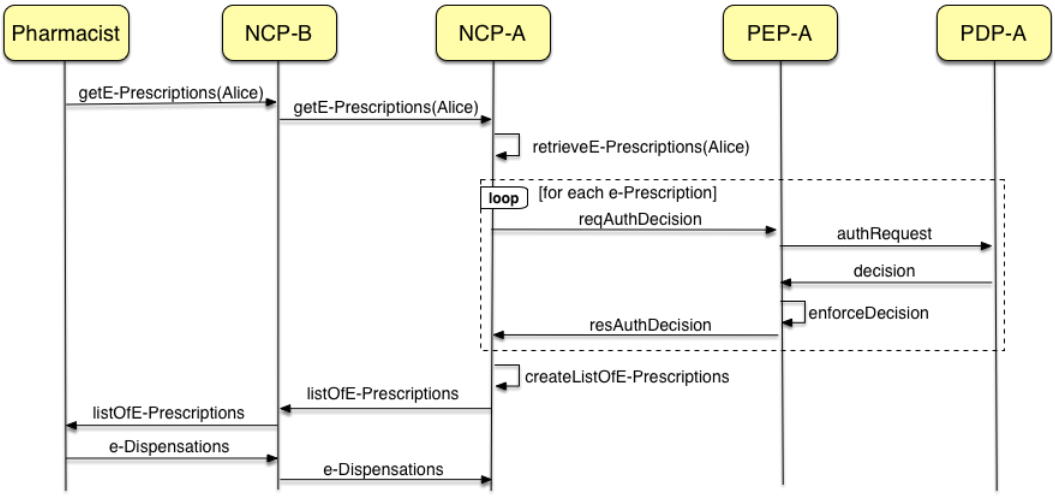

Fig. 2: e-Prescription service protocol

3 An e-Health case study

The case study we consider throughout this paper concerns the provision of e-Health services for exchanging private health data. In this context, we will show that access control policies expressed in FACPL can control accesses to health data in order to preserve data confidentiality and integrity.

The exchange of patients health data among European points of care (such as clinics, hospitals, pharmacies, etc.) has been pursued by the EU through the large scale pilot project epSOS (`http://www.epsos.eu`). The goal is to establish a suite of standardised data exchanging services for facilitating the cross-border interoperability [27] of the EU country healthcare systems and professionals (such as doctors, nurses, pharmacists, etc.), thus ultimately improving the effectiveness of healthcare treatments to EU citizens that are abroad. These services must respect a set of requirements in order to comply with country-specific legislations [16, 44] and to enforce the *patient informed consent*, i.e. the patients informed indications pertaining to personal data processing.

As a case study for this paper, we consider the *electronic prescription* (e-Prescription) service. This service allows EU patients, while staying in a foreign country B participating to the project, to have dispensed a medicine prescribed by a doctor in the country A where the patient is insured. The protocol implemented by this service is illustrated in the message sequence diagram in Figure 2. The e-Prescription service helps pharmacists in country B to retrieve (and properly convert) e-Prescriptions from country A; this is due to trusted actors named National Contact Points (NCPs). Therefore, once a pharmacist has identified the patient (Alice), the remote access is requested to the local NCP (NCP-B), which in its own turn contacts the remote NCP (NCP-A)[3]. The latter one retrieves the e-Prescriptions of the patient from the national infrastructure and, for each e-Prescription, performs through PEP-A an authorisation check against the patient informed consent. In details, PEP-A asks PDP-A to evaluate the pharmacist request with respect to the e-Prescription and the policies expressing the patient consent. Once all decisions are enforced by PEP-A, NCP-A creates the list of e-Prescriptions, by transcoding and translating them into the code system and language of the country B. Finally, the pharmacist dispenses the medicine to the patient and updates the e-Prescription, i.e. it returns e-Dispensation documents.

Starting from the epSOS specifications, we deduced a set of business requirements concerning the e-Prescription service. To streamline the presentation, we explicitly report in Table 1 all and only those requirements authorising some actions. Hence, every action not explicitly authorised is forbidden. For instance, it is not allowed to pharmacists to write e-Prescriptions, which is instead allowed to doctors

[3] For the sake of presentation, we abstract from the authentication process carried out by the pharmacist to ascertain the patient identity.

exhibiting specific permissions. All the requirements are self-explanatory. We just want to point out that the first three requirements deal with access restrictions, while the other ones deal with additional functionalities that sophisticated access control systems, like the one we present, can provide.

Tab. 1: Requirements for the e-Prescription service

#	*Description*
1	Doctors with e-Pre-Read and e-Pre-Write permissions can write e-Prescriptions
2	Doctors with e-Pre-Read permission can read e-Prescriptions
3	Pharmacists with e-Pre-Read permission can read e-Prescriptions
4	Authorised user accesses must be recorded by the system
5	Patients must be informed of unauthorised access attempts
6	Data exchanged should be compressed

4 The FACPL Language

In this section we present FACPL, the language we propose for defining high-level access control policies and requests. First, we introduce its syntax (Section 4.1). Then, we informally explain the semantics of its linguistic constructs (Section 4.2) and employ them to implement the access control system of the e-Health case study (Section 4.3).

4.1 Syntax

Intuitively, FACPL policies are hierarchically structured lists of elements containing controls on the value of attributes that should be provided by FACPL access requests. Together with permit or deny decisions, policies specify the combining algorithms to be used in their evaluation and the obligations for the enforcement process.

Formally, the syntax of FACPL is reported in Table 2. It is given through EBNF-like grammars, where as usual the symbol ? stands for optional items, $*$ for (possibly empty) sequences, and $+$ for non-empty sequences.

A top-level term is a *Policy Authorisation System (PAS)* encompassing the specifications of a PEP and a PDP. The PEP is defined by the *enforcement algorithm* applied for establishing how decisions have to be enforced, e.g. if only decisions permit and deny are admissible, or also not-app and indet can be returned. The PDP is instead defined by a policy, or by a sequence of policies and an algorithm for combining the results of the evaluation of these policies.

A *policy* is made of a sequence of fields separated by keywords. It can be either a basic authorisation *rule* or a *policy set* collecting rules and other policy sets, so that policies can be hierarchically structured. A rule specifies an *effect*, that is the permit or deny decision returned when the rule is successfully evaluated, a *target*, that is an expression indicating the set of access requests to which the rule applies, and a sequence of obligations, that is actions to be discharged by the enforcement process. A policy set specifies a target, a sequence of enclosed policies along with an algorithm for combining the results of their evaluation, and two sequences of obligations, one to be discharged if the resulting effect is permit, the other if it is deny. Obligation sequences may be empty, while policy sequences cannot.

An attribute *name* refers to the literal value associated to the attribute. The name is structured in the form $Identifier/Identifier$, where the first identifier stands for a category name and the second for an attribute name. For example, the structured name subject/role represents the value of the attribute role within the category subject. Categories permit a fine-grained classification of attributes, varying from the usual categories of access control, i.e. *subject*, *resource* and *action*, to possibly application-dependent ones.

Expressions are built from attribute names and *literal* values, i.e. booleans, doubles, strings, and dates, by using standard operators. As usual, string values are written as sequences of characters delimited by double quotes.

Combining algorithms offer different strategies to merge the decisions resulting from the evaluation of various policies (e.g. the p-over$_\delta$ algorithm states that decision permit takes precedence over the others).

Tab. 2: Syntax of FACPL

Policy Auth. Systems	PAS ::=	$\{\,\mathsf{pep} : EnfAlg \quad \mathsf{pdp} : PDP\,\}$
Enforcement algorithms	$EnfAlg$ ::=	base | deny-biased | permit-biased
Policy Decision Points	PDP ::=	$Policy \mid \{Alg \quad \mathsf{policies} : Policy^+\}$
Combining algorithms	Alg ::=	p-over$_\delta$ | d-over$_\delta$ | d-unless-p$_\delta$ | p-unless-d$_\delta$ | first-app$_\delta$ | one-app$_\delta$ | weak-con$_\delta$ | strong-con$_\delta$
Instantiation strategies	δ ::=	greedy | all
Policies	$Policy$ ::=	$Rule$ $\mid \{Alg \quad \mathsf{target} : Expr \quad \mathsf{policies} : Policy^+ \quad \mathsf{obl\text{-}p} : Obligation^* \quad \mathsf{obl\text{-}d} : Obligation^*\,\}$
Rules	$Rule$ ::=	$(Effect \quad \mathsf{target} : Expr \quad \mathsf{obl} : Obligation^*\,)$
Effects	$Effect$ ::=	permit | deny
Obligations	$Obligation$ ::=	$[\,ObType \quad PepAction(Expr^*)\,]$
Obligation types	$ObType$ ::=	m | o
Expressions	$Expr$ ::=	$Name \mid Value \mid \mathsf{and}(Expr, Expr) \mid \mathsf{or}(Expr, Expr) \mid \mathsf{not}(Expr)$ $\mid \mathsf{equal}(Expr, Expr)\mathsf{in}(Expr, Expr) \mid \mathsf{greater\text{-}than}(Expr, Expr) \mid \mathsf{add}(Expr, Expr)$ $\mid \mathsf{subtract}(Expr, Expr) \mid \mathsf{divide}(Expr, Expr) \mid \mathsf{multiply}(Expr, Expr)$
Attribute names	$Name$ ::=	$Identifier/Identifier$
Literal values	$Value$ ::=	true | false | $Double$ | $String$ | $Date$
Requests	$Request$ ::=	$(Name, Value)^+$

Tab. 3: Auxiliary syntax for FACPL responses

PDP responses	$PDPResponse$ ::=	$\langle\, Decision \quad IObligation^*\rangle$
Decisions	$Decision$ ::=	permit | deny | not-app | indet
Instantiated oblig.	$IObligation$ ::=	$[\,ObType \quad PepAction(Value^*)\,]$

They can be specialised by choosing different strategies for the instantiation of obligations (e.g. the greedy strategy states that only the obligations resulting from the actually evaluated policies are returned). In the algorithm names, p and d are shortcuts for permit and deny, respectively.

An *obligation* specifies a type, i.e. mandatory (m) or optional (o), and identifier and arguments of an action to be performed by the PEP. The set of action identifiers accepted by the PEP can be chosen, from time to time, according to the specific application (therefore, *PepAction* is intentionally left unspecified). Action arguments are expressions.

A *request* consists of a (non-empty) sequence of *attributes*, i.e. name-value pairs, that enumerate request credentials in the form of literal values. *Multivalued attributes*, i.e. names associated to a set of values, are rendered as multiple attributes sharing the same name.

The responses resulting from the evaluation of FACPL requests are written using the auxiliary syntax reported in Table 3. The two-stage evaluation process described in Section 2 produces two different kinds of responses: *PDP responses* and *decisions* (i.e. responses by the PEP). The former ones, in case of decision permit and deny, pair the decision with a (possibly empty) sequence of instantiated obligations. An *instantiated obligation* is a pair made of a type (i.e., m or o) and an action whose arguments are values.

To simplify notations, in the sequel we will omit the keyword preceding a sub-term generated by the grammar in Table 2 whenever the sub-term is missing or is the expression true. Thus, e.g., the rule (deny target : true obl :) will be simply written as (deny). Moreover, when in the *PDPResponse* the sequence of instantiated obligations is empty, we sometimes write *Decision* instead of $\langle Decision\rangle$.

4.2 Informal Semantics

We now informally explain how the FACPL linguistic constructs are dealt with in the evaluation process of access requests described in Section 2. We first present the PDP authorisation process and then the PEP enforcement process.

When the PDP receives an access request, first it evaluates the request on the basis of the available policies. Then, it determines the resulting decision by combining the decisions returned by these policies through the top-level combining algorithm.

The evaluation of a policy with respect to a request starts by checking its applicability to the request, which is done by evaluating the expression defining its target. Let us suppose that the applicability holds, i.e. the expression evaluates to true. In case of rules, the rule effect is returned. In case of policy sets, the result is obtained by evaluating the contained policies and combining their evaluation results through the specified algorithm. In both cases, the evaluation ends with the instantiation of the enclosed obligations. Let us suppose now that the applicability does not hold. If the expression evaluates to false, the policy evaluation returns not-app, while if the expression returns an error or a non-boolean value, the policy evaluation returns indet. Clearly, the target of enclosed policies may refine that of the enclosing ones, while a policy with target expression true (resp., false) applies to all (resp., no) requests.

Evaluating expressions amounts to apply operators and to *resolve* the attribute names occurring within, that is to determine the value corresponding to each such name. This value can either be contained in the request or retrieved from the environment by the context handler (steps 5-8 in Figure 1). Thus, if an attribute with that name is missing in the request and its retrieval by the context handler fails, the special value $\bot$ is returned. Taking the value $\bot$ apart from errors permits both carefully managing those requests only containing a limited set of attributes and reasoning on the role of missing attributes in policy evaluation (see Section 7 for details).

It is worth noticing that the syntax of policies, and in particular that of attribute names and expressions, does not consider types. Indeed, we want a policy to provide a response to any request, not only to those complying with the expected type of (the values referred by) the attribute names controlled by the policy. Since we do not filter requests on the base of the type of their attributes, we cannot in general statically ensure that expressions within policies are well-typed. Consequently, errors will be generated at evaluation-time , and possibly managed, when expression operators are applied to arguments of unexpected type.

Indeed, the evaluation of expressions takes into account the types of the operators' arguments, and possibly returns the special values $\bot$ and error. In details, if the arguments are of the expected type, the operator is applied, else, i.e. at least one argument is error, error is returned; otherwise, i.e. at least one argument is $\bot$ and none is error, $\bot$ is returned. The operators and and or implement a different treatment of these special values. Specifically, and returns true if both operands are true, false if at least one operand is false, $\bot$ if at least one operand is $\bot$ and none is false or error, and error otherwise (e.g. when an operand is not a boolean value). The operator or is the dual of and. Hence, and and or may mask $\bot$ and error. Instead, the unary operator not only swaps values true and false and leaves $\bot$ and error unchanged. In the rest, we use operators and and or in infix notation, and assume that they are commutative and associative, and that operator and takes precedence over or.

The evaluation of a rule ends with the instantiation of all the enclosed obligations, while that of a policy set ends with the instantiation of all the obligations in the sequence corresponding to the decision calculated for the policy. The instantiation of an obligation consists in evaluating every expression argument of the enclosed action. If an error occurs, the policy decision is changed to indet. Otherwise, the instantiated obligations are paired with the policy decision to form the PDP response.

Evaluating a policy set requires the application of the specified algorithm for combining the decisions resulting from the evaluation of various policies and, thus, resolving possible conflicts, e.g. whenever both decisions permit and deny occur. Given a sequence of policies in input, the combining algorithms prescribe the sequential evaluation of the given policies and behave as follows:

- p-over$_\delta$ (d-over$_\delta$ is specular): if the evaluation of a policy returns permit, then the result is permit. In other words, permit takes precedence, regardless of the result of any other policy. Instead, if at least

one policy returns deny and all others return not-app or deny, then the result is deny. If all policies return not-app, then the result is not-app. In the remaining cases, the result is indet.

- d-unless-p$_\delta$ (p-unless-d$_\delta$ is specular): similarly to p-over$_\delta$, this algorithm gives precedence to permit over deny, but it always returns deny in all the other cases.
- first-app$_\delta$: the algorithm returns the evaluation result of the first policy in the sequence that does not return not-app, otherwise the result is not-app.
- one-app$_\delta$: when exactly one policy is applicable, the result of the algorithm is that of the applicable policy. If no policy applies, the algorithm returns not-app, while if more than one policy is applicable, it returns indet.
- weak-con$_\delta$: the algorithm returns permit (resp., deny) if some policies return permit (resp., deny) and no other policy returns deny (resp., permit); if both decisions are returned, the algorithm returns indet. If policies only return not-app or indet, then indet, if present, prevails.
- strong-con$_\delta$: this algorithm is the stronger version of the previous one, in the sense that to obtain permit (resp., deny) all policies have to return permit (resp., deny), otherwise indet is returned. If all policies return not-app then the result is not-app.

The algorithms described in the first four items above are those popularised by XACML. They combine decisions either according to a given precedence criterium or to policy applicability. The remaining two algorithms, instead, are borrowed from [31] and compute the combined decision by achieving different forms of consensus.

If the resulting decision is permit or deny, each algorithm also returns the sequence of instantiated obligations according to the chosen instantiation strategy δ. There are two possible strategies. The all strategy requires evaluation of all policies in the input sequence and returns the instantiated obligations pertaining to all decisions. Instead, the greedy strategy prescribes that, as soon as a decision is obtained that cannot change due to evaluation of subsequent policies in the input sequence, the execution halts. Hence, the result will not consider the possibly remaining policies and only contains the obligations already instantiated. Therefore, the instantiation strategies mainly affect the amount of instantiated obligations possibly returned. The greedy strategy, that reflects the management of obligations in XACML, may significantly improve the policy evaluation performance. Instead, the all strategy may require additional workload but, on the other hand, ensures that all the policies and their obligations are always taken into account.

As last step, the calculated PDP response is sent to the PEP for the enforcement. To this aim, the PEP must discharge all obligations and decide, by means of the chosen enforcement algorithm, how to enforce decisions not-app and indet. The algorithms are those popularised by XACML and, in particular, the deny-biased (resp., permit-biased) one enforces permit (resp., deny) only when all the corresponding obligations are correctly discharged, while enforces deny (resp., permit) in all other cases. Instead, the base algorithm leaves all decisions unchanged but, in case of decisions permit and deny, enforces indet if an error occurs while discharging obligations. This means that obligations not only affect the authorisation process due to their instantiation, but also the enforcement one. However, errors caused by optional obligations, i.e. with type o, are safely ignored.

4.3 Policies for the e-Health case study

We now use the FACPL linguistic abstractions to formalise the requirements for the e-Health case study reported in Table 1. These rules are meant to prevent unauthorised access to patient data and hence to ensure their confidentiality and integrity. The specification of this access control system is introduced bottom-up, from single rules to whole policies, thus illustrating in a step-by-step fashion the combination strategies that could be pursued and their effects.

The system resources to protect via the access control system are *e-Prescriptions*. The access control rules need to deal with requester credentials, i.e. doctor and pharmacist roles, along with their assigned permissions, and with read or write actions.

Requirement (1), allowing doctors to write e-Prescriptions, can be formalised as a *positive* FACPL rule (i.e., a rule with effect permit) as follows

```
( permit  target : equal(subject/role, "doctor")
            and equal(action/id, "write")
            and in("e-Pre-Write", subject/permission)
            and in("e-Pre-Read", subject/permission))
```

The rule target[4] checks if the requester role is doctor, if the action is write, and if the subject's permissions include those for writing and reading an e-Prescription. The control that the resource type is equal to e-Prescription will be performed by the target of the policy enclosing the rule. This, because of the hierarchical processing of FACPL elements, is enough to ensure that the rule will only be applied to e-Prescriptions.

Requirement (2) can be expressed like the previous: it differs for the action identifier and for the required permissions, i.e. only e-Pre-Read. Requirement (3) only differs from the second for the role value.

These three rules, modelling Requirements (1), (2) and (3), can be combined together in a policy set whose target specifies the check on the resource type e-Prescription (again, to improve code readability, we use textual encoding for resources). Since all granted requests are explicitly authorised, choosing the $\text{p-over}_{\text{all}}$ algorithm as combining strategy seems a natural choice. Let thus Policy (P1) be defined as follows

```
{ p-over_all
  target : equal(resource/type, "e-Prescription")
  policies :
     ( permit  target : equal(subject/role, "doctor")
          and equal(action/id, "write")
          and in("e-Pre-Write", subject/permission)
          and in("e-Pre-Read", subject/permission))
     ( permit  target : equal(subject/role, "doctor")
          and equal(action/id, "read")
          and in("e-Pre-Read", subject/permission))
     ( permit  target : equal(subject/role, "pharmacist")
          and equal(action/id, "read")
          and in("e-Pre-Read", subject/permission))
  obl-p : [ m  log(system/time, resource/type, subject/id, action/id) ] }
```
(P1)

Policy (P1) reports not only access controls but also an obligation formalising Assumption (4) about the logging of each authorised access. The arguments of the obligation action are separated by commas to increase their readability.

Let us now consider a FACPL request and evaluate it with respect to Policy (P1). For the sake of presentation, hereafter we write $A \triangleq t$ to assign the symbolic name A to the term t. Let us suppose that doctor Dr. House wants to write an e-Prescription; the corresponding request is defined as follows

```
req1 ≜ (subject/id, "Dr. House")
       (subject/role, "doctor") (action/id, "write")
       (resource/type, "e-Prescription")
       (subject/permission, "e-Pre-Read")
       (subject/permission, "e-Pre-Write")  ...
```

[4] To improve code readability, we use the infix operators, a textual notation for permissions and an additional check on the subject role. Of course, in a setting with semantically different roles, a standardised permission-based coding, e.g. HL7 (`http://www.hl7.org`), should be used for defining role checks.

where attributes are organised into the categories *subject*, *resource* and *action*. Additional attributes possibly included in the request are omitted because they are not relevant for this evaluation. Notice that subject/permission is a multivalued attribute and it is properly handled in the previous rules by using the in operator, which verifies the membership of its first argument to the set that forms its second argument.

The authorisation process of req1 returns a permit decision. In fact, the request matches the policy target, as the resource type is e-Prescription, and exposes all the permissions required in the first rule for the write action and the doctor role. The response, that is a permit including a log obligation, is defined, e.g., as follows

$$\langle\ \mathsf{permit}\ \ [\,\mathsf{m}\ \ \mathsf{log}(\text{2016-10-22 10:15:12}, \text{“e-Prescription”}, \text{“Dr. House”}, \text{“write”})\,]\rangle$$

The instantiated obligation indicates that the PDP succeeded in retrieving and evaluating all the attributes occurring within the arguments of the action; run-time information, such as the current time, is retrieved through the context handler.

The evaluation of req1 returns the expected result. We might be led to believe that due to the simplicity of Policy (P1), this is true for all requests. However, this correctness property cannot be taken for granted as, in general, even though the meaning of a rule is straightforward, this may not be the case for a combination of rules. Depending on the chosen combination strategy, some unexpected results can arise. For example, a request by a pharmacist for a write action on an e-Prescription is not explicitly allowed by the requirements in Table 1; hence, it should be forbidden. However, the corresponding request

$$\begin{aligned}\mathsf{req2} \triangleq\ & (\mathsf{subject/id}, \text{“Dr. Wilson”})\\ & (\mathsf{subject/role}, \text{“pharmacist”})\,(\mathsf{action/id}, \text{“write”})\\ & (\mathsf{resource/type}, \text{“e-Prescription”})\\ & (\mathsf{subject/permission}, \text{“e-Pre-Read”})\quad \ldots\end{aligned}$$

would evaluate to not-app. In fact, all enclosed rules do not apply (i.e., their targets do not match) and the resulting not-app decisions are combined by the $\mathsf{p\text{-}over}_{\mathsf{all}}$ algorithm to not-app as well. Therefore, the enforcement algorithm of the PEP is entrusted with the task of taking the final decision for request req2. Even though this is correct in a setting where the PEP is well-defined, e.g. the epSOS system, it is not recommended when design assumptions on the PEP implementation are missing. In fact, a biased algorithm might transform not-app into permit, possibly causing unauthorised accesses.

To prevent not-app decisions to be returned by the policy, we can replace the combining algorithm of Policy (P1) with the $\mathsf{d\text{-}unless\text{-}p}_{\mathsf{all}}$ one. This implies that deny is taken as the default decision and is returned whenever no rule returns permit. Alternatively, we can get the same achievement by using a policy set defined as the combination, through the $\mathsf{p\text{-}over}_{\mathsf{all}}$ algorithm, of Policy (P1) and a rule forbidding all accesses. This rule is simply defined as (deny): the absence of the target and the *negative* effect means that it always returns deny. Now, let Policy (P2) be defined as

$$\begin{aligned}&\{\,\mathsf{p\text{-}over}_{\mathsf{all}}\\ &\quad \mathsf{policies}: \{\ldots \mathit{Policy}\ (\mathit{P1})\ldots\}\ \ (\mathsf{deny})\\ &\quad \mathsf{obl\text{-}p}: [\,\mathsf{o}\ \ \mathsf{compress}(\,)\,]\\ &\quad \mathsf{obl\text{-}d}: [\,\mathsf{m}\ \ \mathsf{mailTo}(\mathsf{resource/patient\text{-}mail},\\ &\qquad\qquad \text{“Data request by unauthorised subject”})\,]\,\}\end{aligned}\tag{P2}$$

Policy (P2) reports two obligations formalising, respectively, the last two requirements of Table 1: (i) a patient is informed about unauthorised attempts to access her data by means of an obligation for the effect deny and (ii) if possible, data are exchanged in compressed form by means of an obligation for the effect permit. Notably, the type ‘optional’ is exploited so that compressed exchanges are not strictly required but, e.g., only whenever the corresponding service is available.

Policy (P2) can be used as a basis for the definition of the *patient informed consent* (see Section 3). For instance, Alice's policy for the management of her health data could be simply obtained by adding a check on the patient identifier to which the policy applies, such as $\mathsf{target}: \mathsf{equal}(\text{“Alice”}, \mathsf{resource/patient\text{-}id})$, to Policy (P2). In this way, Alice grants access to her e-Prescription data to the healthcare professionals

Tab. 4: Correspondence between syntactic and semantic domains

Syntactic category	**Generic synt. elem.**	**Semantic function**	**Syntactic domain**	**Semantic domain**
Attribute names	n		*Name*	
Literal values	v		*Value*	
Requests	*req*	$\mathcal{R}$	*Request*	$R \triangleq \mathit{Name} \to (\mathit{Value} \cup 2^{\mathit{Value}} \cup \{\bot\})$
Expressions	*expr*	$\mathcal{E}$	*Expr*	$R \to \mathit{Value} \cup 2^{\mathit{Value}} \cup \{\mathsf{error}, \bot\}$
Effects	e		*Effect*	
Obligation Types	t		*ObType*	
Pep Actions	*pepAct*		*PepAction*	
Instantiated obligations	*io*		*IObligation*	
Obligations	o	$\mathcal{O}$	*Obligation*	$R \to \mathit{IObligation} \cup \{\mathsf{error}\}$
PDP Responses	*res*		*PDPReponse*	
Policies	p	$\mathcal{P}$	*Policy*	$R \to \mathit{PDPReponse}$
Policy Decision Points	*pdp*	$\mathcal{P}dp$	*PDP*	$R \to \mathit{PDPReponse}$
Combining algorithms	a	$\mathcal{A}$	$\mathit{Alg} \times \mathit{Policy}^+$	$R \to \mathit{PDPReponse}$
Decisions	*dec*		*Decision*	
Enforcement algorithms	*ea*	$\mathcal{E}A$	*EnfAlg*	$\mathit{PDPReponse} \to \mathit{Decision}$
Policy Auth. System	*pas*	$\mathcal{P}as$	*PAS*	$\mathit{Request} \to \mathit{Decision}$

that satisfy the requirements expressed in her consent policy. Another patient expressing a more restrictive consent, where e.g. writing of e-Prescriptions is disabled, will have a similar policy set where the rule modelling Requirement (1) is not included. In a more general perspective, the PDP could have a policy set for each patient, that encloses the policies expressing the consent explicitly signed by the patient. This is the approach followed, e.g., in the Austrian e-Health platform (`http://www.elga.gv.at/`).

As shown before, it could be challenging to identify unexpected authorisations and to determine whether policy fixes affect authorisations that should not be altered. The combination of a large number of complex policies is indeed an error-prone task that has to be supported with effective analysis techniques. Therefore we equip FACPL with a formal semantics and then define a constraint-based analysis providing effective supporting techniques for the verification of properties on policies.

5 FACPL Formal Semantics

In this section, we present the formal semantics of FACPL by formalising the evaluation process introduced in Section 2 and detailed in Section 4.2. The semantics is defined by following a denotational approach which means that

- we introduce some semantic functions mapping each FACPL syntactic construct to an appropriate *denotation*, that is an element of a semantic domain representing the meaning of the construct;
- the semantic functions are defined in a *compositional* way, so that the semantic of each construct is formulated in terms of the semantics of its sub-constructs.

To this purpose, we specify a family of semantic functions mapping each syntactic domain to a specific semantic domain. These functions are inductively defined on the FACPL syntax through appropriate semantic clauses following a 'point-wise' style. For instance, on the syntactic domain *Policy* representing all FACPL policies, we formalise the function $\mathcal{P}$ that defines a semantic domain mapping FACPL requests to PDP responses.

In the sequel, we convene that the application of the semantic functions is left-associative, omits parenthesis whenever possible, and surrounds syntactic objects with the emphatic brackets $[\![$ and $]\!]$ to increase readability. For instance, $\mathcal{E}[\![n]\!]r$ stands for $(\mathcal{E}(n))(r)$ and indicates the application of the semantic function

$\mathcal{E}$ to (the syntactic object) n and (the semantic object) r. We also assume that each nonterminal symbol in Tables 2 and 3 (defining the FACPL syntax) denominates the set of constructs of the syntactic category defined by the corresponding EBNF rule, e.g. the nonterminal *Policy* identifies the set of all FACPL policies. The used notations are summarised in Table 4 (the missing semantic domains coincide with the corresponding syntactic ones).

In the rest of this section, we detail the semantics of requests (Section 5.1), PDP (Sections 5.2 and 5.3), PEP (Section 5.4), Policy Authorisation System (Section 5.5) and present some properties of the semantics (Section 5.6).

5.1 Semantics of Requests

The meaning of a request[5] is an element of the set $R \triangleq \mathit{Name} \to (\mathit{Value} \cup 2^{\mathit{Value}} \cup \{\bot\})$, that is a total function that maps attribute names to either a literal value, or a set of values (in case of multivalued attributes), or the special value $\bot$ (if the value for an attribute name is missing). The mapping from a request to its meaning is given by the semantic function $\mathcal{R} : \mathit{Request} \to R$, defined as follows:

$$\mathcal{R}[\![(n', v')]\!]n = \begin{cases} v' \texttt{ if } n = n' \\ \bot \texttt{ otherwise} \end{cases}$$
$$\mathcal{R}[\![(n_i, v_i)^+(n', v')]\!]n = \begin{cases} \mathcal{R}[\![(n_i, v_i)^+]\!]n \Cup v' \texttt{ if } n = n' \\ \mathcal{R}[\![(n_i, v_i)^+]\!]n \quad \texttt{otherwise} \end{cases} \tag{S-1}$$

The semantics of a request, which is a function $r \in R$, is thus inductively defined on the length of the request. To deal with multivalued attributes we introduce the operator $\Cup$, which is straightforwardly defined by case analysis on the first argument as follows

$$v \Cup v' = \{v, v'\} \qquad V \Cup v' = V \cup \{v'\} \qquad \bot \Cup v' = v'$$

where we let $V \in 2^{\mathit{Value}}$.

5.2 Semantics of the Policy Decision Process

We start defining the semantics of expressions and obligations that will be then exploited for defining the semantics of policies.

In Table 5 we report (an excerpt of) the clauses defining the function $\mathcal{E} : \mathit{Expr} \to (R \to \mathit{Value} \cup 2^{\mathit{Value}} \cup \{\mathsf{error}, \bot\})$ modelling the semantics of expressions. This means that the semantics of an expression is a function of the form $R \to \mathit{Value} \cup 2^{\mathit{Value}} \cup \{\mathsf{error}, \bot\}$ that, given a request, returns a literal value, or a set of values, or the special value $\bot$, or an error (e.g. when an argument of an operator has unexpected type). The evaluation order of sub-expressions is not relevant, as they do not generate side-effects.

The first raw of the table contains the clauses for basic expressions, i.e. attribute names and literal values. The semantics of the expression formed by a name n is a function that, given a (semantic) request r in input, returns the value that r associates to n. This is written as the clause $\mathcal{E}[\![n]\!]r = r(n)$. Similarly, the case of a value v is a function that always returns the value itself, that is the clause $\mathcal{E}[\![v]\!]r = v$.

The remaining clauses, one for each operator, present (an excerpt of) the semantics of expression operators. In particular, each clause uses straightforward semantic operators for composing denotations (e.g. $=$ corresponds to equal), and implements the management strategy for the special values $\bot$ and error. The clauses establish that error takes precedence over $\bot$ and is returned every time the operator arguments have unexpected types; whereas $\bot$ is returned when at least an argument is $\bot$ and there is no error. The clauses of operators and and or possibly mask these special values by implementing the behaviour informally described

[5] For simplicity sake, here we assume that, when the evaluation of a request takes place, the original request has been already enriched with the information that would be retrieved at run-time from the environment by the context handler (steps 5-8 in Figure 1).

Tab. 5: Semantics of (an excerpt of) FACPL Expressions (T stands for one of the sets of literal values or for the powerset of the set of all literal values, and $i, j \in \{1, 2\}$ with $i \neq j$)

$$\mathcal{E}[\![n]\!]r = r(n)$$

$$\mathcal{E}[\![v]\!]r = v$$

$$\mathcal{E}[\![\mathsf{or}(expr_1, expr_2)]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{E}[\![expr_1]\!]r = \mathsf{true} \vee \mathcal{E}[\![expr_2]\!]r = \mathsf{true} \\ \mathsf{false} & \texttt{if } \mathcal{E}[\![expr_1]\!]r = \mathcal{E}[\![expr_2]\!]r = \mathsf{false} \\ \bot & \texttt{if } \mathcal{E}[\![expr_i]\!]r = \bot \wedge \mathcal{E}[\![expr_j]\!]r \in \{\mathsf{false}, \bot\} \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

$$\mathcal{E}[\![\mathsf{and}(expr_1, expr_2)]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{E}[\![expr_1]\!]r = \mathcal{E}[\![expr_2]\!]r = \mathsf{true} \\ \mathsf{false} & \texttt{if } \mathcal{E}[\![expr_1]\!]r = \mathsf{false} \vee \mathcal{E}[\![expr_2]\!]r = \mathsf{false} \\ \bot & \texttt{if } \mathcal{E}[\![expr_i]\!]r = \bot \wedge \mathcal{E}[\![expr_j]\!]r \in \{\mathsf{true}, \bot\} \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

$$\mathcal{E}[\![\mathsf{not}(expr)]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{E}[\![expr]\!]r = \mathsf{false} \\ \mathsf{false} & \texttt{if } \mathcal{E}[\![expr]\!]r = \mathsf{true} \\ \bot & \texttt{if } \mathcal{E}[\![expr]\!]r = \bot \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

$$\mathcal{E}[\![\mathsf{add}(expr_1, expr_2)]\!]r = \begin{cases} (\mathcal{E}[\![expr_1]\!]r + \mathcal{E}[\![expr_2]\!]r) & \texttt{if } \mathcal{E}[\![expr_1]\!]r, \mathcal{E}[\![expr_2]\!]r \in \mathit{Double} \\ \bot & \texttt{if } \mathcal{E}[\![expr_i]\!]r = \bot \ \wedge \mathcal{E}[\![expr_j]\!]r \neq \mathsf{error} \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

$$\mathcal{E}[\![\mathsf{in}(expr_1, expr_2)]\!]r = \begin{cases} (\mathcal{E}[\![expr_1]\!]r \in \mathcal{E}[\![expr_2]\!]r) & \texttt{if } \mathcal{E}[\![expr_1]\!]r \in T \ \wedge \ \mathcal{E}[\![expr_2]\!]r \in 2^T \\ \bot & \texttt{if } \mathcal{E}[\![expr_i]\!]r = \bot \ \wedge \mathcal{E}[\![expr_j]\!]r \neq \mathsf{error} \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

$$\mathcal{E}[\![\mathsf{equal}(expr_1, expr_2)]\!]r = \begin{cases} (\mathcal{E}[\![expr_1]\!]r = \mathcal{E}[\![expr_2]\!]r) & \texttt{if } \mathcal{E}[\![expr_1]\!]r, \mathcal{E}[\![expr_2]\!]r \in T \\ \bot & \texttt{if } \mathcal{E}[\![expr_i]\!]r = \bot \\ & \quad \wedge \mathcal{E}[\![expr_j]\!]r \neq \mathsf{error} \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

in Section 4.2. It is worth noticing that the explicit management of missing attributes and evaluation errors ensures a full account of crucial aspects of access control policy evaluation, usually neglected by other proposals from the literature (see, e.g., [24, 40, 2]). The only proposals considering the role of missing attributes are those in [8, 10], but they only consider a simplified policy language and assume that expressions cannot generate errors.

Function $\mathcal{E}$ is straightforwardly extended to sequences of expressions by the following clauses

$$\begin{aligned} \mathcal{E}[\![\epsilon]\!]r &= \epsilon \\ \mathcal{E}[\![expr' \ \ expr^*]\!]r &= \mathcal{E}[\![expr']\!]r \bullet \mathcal{E}[\![expr^*]\!]r \end{aligned} \tag{S-2}$$

The operator $\bullet$ denotes concatenation of sequences of semantic elements and ϵ denotes the empty sequence. We assume that $\bullet$ is strict on error and $\bot$, i.e. error is returned whenever an error or $\bot$ is in the sequence. Therefore, the evaluation of $\mathcal{E}[\![expr^*]\!]r$ fails if any of the expressions forming $expr^*$ evaluates to error or $\bot$.

The semantics of the instantiation of obligations is formalised by the function $\mathcal{O} : \ \mathit{Obligation} \rightarrow \ (R \rightarrow \mathit{IOblgation} \cup \{\mathsf{error}\})$ defined by the clause

$$\mathcal{O}[\![[\, t \ \ pepAct(expr^*)\,]]\!]r = \begin{cases} [\, t \ \ pepAct(w^*)\,] & \texttt{if } \ \mathcal{E}[\![expr^*]\!]r = w^* \\ \mathsf{error} & \texttt{otherwise} \end{cases} \tag{S-3a}$$

where w stands for a literal value or a set of literal values. Thus, given a request, the instantiation of an obligation returns an instantiated obligation, if the evaluation of every expression argument of the action returns a value. Otherwise, it returns an error.

Function $\mathcal{O}$ is straightforwardly extended to sequences of obligations as follows

$$\mathcal{O}[\![\epsilon]\!]r = \epsilon \qquad \mathcal{O}[\![o' o^*]\!]r = \mathcal{O}[\![o']\!]r \bullet \mathcal{O}[\![o^*]\!]r \tag{S-3b}$$

Notably, a sequence of instantiated obligations is returned only if every obligation in the sequence is successfully instantiated; otherwise, error is returned (indeed, $\bullet$ is strict on error).

We can now define the semantics of a policy as a function that, given a request, returns an authorisation decision paired with a (possibly empty) sequence of instantiated obligations. Formally, it is given by the function $\mathcal{P} : \mathit{Policy} \to (R \to \mathit{PDPReponse})$ that has two defining clauses: one for rules and one for policy sets. The clause for rules is

$$\mathcal{P}[\![(e \;\; \mathsf{target} : \mathit{expr} \;\; \mathsf{obl} : o^* \,)]\!]r = \begin{cases} \langle e \;\; io^* \rangle & \texttt{if } \mathcal{E}[\![\mathit{expr}]\!]r = \mathsf{true} \\ & \qquad \wedge\, \mathcal{O}[\![o^*]\!]r = io^* \\ \mathsf{not\text{-}app} & \texttt{if } \mathcal{E}[\![\mathit{expr}]\!]r = \mathsf{false} \\ & \qquad \vee\, \mathcal{E}[\![\mathit{expr}]\!]r = \bot \\ \mathsf{indet} & \texttt{otherwise} \end{cases} \tag{S-4a}$$

Thus, the rule effect is returned as a decision when the target evaluates to true, which means that the rule applies to the request, and all obligations are successfully instantiated. In this case, the instantiated obligations are also part of the response. Otherwise, it could be the case that *(i)* the rule does not apply to the request, i.e. the target evaluates to false or to $\bot$, or that *(ii)* an error has occurred while evaluating the target or instantiating the obligations.

The semantics of policy sets relies on the semantics of combining algorithms. Indeed, as detailed in Section 5.3, we use a semantic function $\mathcal{A}$ to map each combining algorithm a to a function that, to a sequence of policies, associates a function from requests to PDP responses. The clause for policy sets is

$$\mathcal{P}[\![\{a \;\; \mathsf{target} : \mathit{expr} \;\; \mathsf{policies} : p^+ \;\; \mathsf{obl\text{-}p} : o_p^* \;\; \mathsf{obl\text{-}d} : o_d^* \,\}]\!]r = \begin{cases} \langle \mathsf{permit} \;\; io_1^* \bullet io_2^* \rangle & \texttt{if } \mathcal{E}[\![\mathit{expr}]\!]r = \mathsf{true} \\ & \qquad \wedge\, \mathcal{A}[\![a, p^+]\!]r = \langle \mathsf{permit} \;\; io_1^* \rangle \\ & \qquad \wedge\, \mathcal{O}[\![o_p^*]\!]r = io_2^* \\ \langle \mathsf{deny} \;\; io_1^* \bullet io_2^* \rangle & \texttt{if } \mathcal{E}[\![\mathit{expr}]\!]r = \mathsf{true} \\ & \qquad \wedge\, \mathcal{A}[\![a, p^+]\!]r = \langle \mathsf{deny} \;\; io_1^* \rangle \\ & \qquad \wedge\, \mathcal{O}[\![o_d^*]\!]r = io_2^* \\ \mathsf{not\text{-}app} & \texttt{if } \mathcal{E}[\![\mathit{expr}]\!]r = \mathsf{false} \\ & \qquad \vee\, \mathcal{E}[\![\mathit{expr}]\!]r = \bot \\ & \qquad \vee\, (\mathcal{E}[\![\mathit{expr}]\!]r = \mathsf{true} \\ & \qquad\quad \wedge\, \mathcal{A}[\![a, p^+]\!]r = \mathsf{not\text{-}app}) \\ \mathsf{indet} & \texttt{otherwise} \end{cases} \tag{S-4b}$$

Thus, the policy set applies to the request when the target evaluates to true, the semantic of the combining algorithm a (which is applied to the enclosed sequence of policies and the request) returns the effect e and a sequence of instantiated obligations io_1^*, and all the enclosed obligations for the effect e are successfully instantiated and return a sequence io_2^*. In this case, the PDP response contains e and the concatenation of the sequences io_1^* and io_2^*. Instead, if the target evaluates to false or to $\bot$, or the combining algorithm returns not-app, the policy set does not apply to the request. The response is indet in the remaining cases, i.e. when an error occurred in the evaluation of the target or of the obligations, or when the evaluation of the combining algorithm returned indet.

Finally, the semantic of a PDP is that function from requests to PDP responses obtained by applying the combining algorithm to the enclosed sequence of policies, i.e.

$$\mathcal{P}dp[\![\{a \;\; \mathsf{policies} : p^+\}]\!]r = \mathcal{A}[\![a, p^+]\!]r \tag{S-5}$$

5.3 Semantics of Combining Algorithms

The semantics of combining algorithms is defined in terms of a family of binary operators. Let alg denote the name of a combining algorithm (i.e., p-over, d-over, etc.); the corresponding semantic operator is identified

Tab. 6: Auxiliary definitions for the semantics of combining algorithms: (a) combination matrix for the $\otimes$p-over operator (res_1 and res_2 indicate the first and the second argument, respectively); (b) definition of the $isFinal_{\mathsf{alg}}(res)$ predicate

(a)

$res_1 \backslash^{res_2}$	$\langle$permit $io_2^*\rangle$	$\langle$deny $io_2^*\rangle$	not-app	indet
$\langle$permit $io_1^*\rangle$	$\langle$permit $io_1^* \bullet io_2^*\rangle$	$\langle$permit $io_1^*\rangle$	$\langle$permit $io_1^*\rangle$	$\langle$permit $io_1^*\rangle$
$\langle$deny $io_1^*\rangle$	$\langle$permit $io_2^*\rangle$	$\langle$deny $io_1^* \bullet io_2^*\rangle$	$\langle$deny $io_1^*\rangle$	indet
not-app	$\langle$permit $io_2^*\rangle$	$\langle$deny $io_2^*\rangle$	not-app	indet
indet	$\langle$permit $io_2^*\rangle$	indet	indet	indet

(b)

$$isFinal_{\mathsf{p\text{-}over}}(res) = \begin{cases} \mathsf{true} & \texttt{if } res.dec = \mathsf{permit} \\ \mathsf{false} & \texttt{otherwise} \end{cases}$$

$$isFinal_{\mathsf{d\text{-}over}}(res) = \begin{cases} \mathsf{true} & \texttt{if } res.dec = \mathsf{deny} \\ \mathsf{false} & \texttt{otherwise} \end{cases}$$

$$isFinal_{\mathsf{d\text{-}unless\text{-}p}}(res) = \begin{cases} \mathsf{true} & \texttt{if } res.dec = \mathsf{permit} \\ \mathsf{false} & \texttt{otherwise} \end{cases}$$

$$isFinal_{\mathsf{p\text{-}unless\text{-}d}}(res) = \begin{cases} \mathsf{true} & \texttt{if } res.dec = \mathsf{deny} \\ \mathsf{false} & \texttt{otherwise} \end{cases}$$

$$isFinal_{\mathsf{first\text{-}app}}(res) = \begin{cases} \mathsf{false} & \texttt{if } res.dec = \mathsf{not\text{-}app} \\ \mathsf{true} & \texttt{otherwsise} \end{cases}$$

$$isFinal_{\mathsf{one\text{-}app}}(res) = \begin{cases} \mathsf{true} & \texttt{if } res.dec = \mathsf{indet} \\ \mathsf{false} & \texttt{otherwsise} \end{cases}$$

$$isFinal_{\mathsf{weak\text{-}con}}(res) = \begin{cases} \mathsf{true} & \texttt{if } res.dec = \mathsf{indet} \\ \mathsf{false} & \texttt{otherwsise} \end{cases}$$

$$isFinal_{\mathsf{strong\text{-}con}}(res) = \begin{cases} \mathsf{true} & \texttt{if } res.dec = \mathsf{indet} \\ \mathsf{false} & \texttt{otherwsise} \end{cases}$$

as $\otimes$alg and is defined by means of a two-dimensional matrix that, given two PDP responses, calculates the resulting combined response. For instance, Table 6(a) reports the combination matrix for the $\otimes$p-over operator. Basically, the matrix specifies the precedences among the permit, deny, not-app and indet decisions, and shows how the resulting (sequence of) instantiated obligations is obtained, i.e. by concatenating the instantiated obligations of the responses whose decision matches the combined one. All other combining algorithms described in Section 4.2, and possibly many others, can be defined in the same manner (see Appendix A).

The semantics of the combining algorithms can be now formalised by the function $\mathcal{A} : Alg \times Policy^+ \rightarrow (R \rightarrow PDPReponse)$. This function is defined in terms of the iterative application of the binary combining operators by means of two definition clauses according to the adopted instantiation strategy: the all strategy always requires evaluation of all policies, while the greedy strategy halts the evaluation as soon as a final decision is determined (i.e. without necessarily taking into account all policies in the sequence). If the all strategy is adopted, the definition clause is as follows

$$\begin{aligned} &\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, p_1 \ \ldots \ p_s]\!]r = \\ &\otimes\mathsf{alg}(\otimes\mathsf{alg}(\ldots \otimes \mathsf{alg}(\mathcal{P}[\![p_1]\!]r, \mathcal{P}[\![p_2]\!]r), \ldots), \mathcal{P}[\![p_s]\!]r) \end{aligned} \tag{S-6a}$$

meaning that the combining operator is sequentially applied to the denotations of all input policies[6]. Instead, if the greedy strategy is used, the definition clause is as follows

$$\begin{aligned} &\mathcal{A}[\![\mathsf{alg}_{\mathsf{greedy}}, p_1 \ \ldots \ p_s]\!]r = \\ &\begin{cases} res_1 & \texttt{if } \mathcal{P}[\![p_1]\!]r = res_1 \ \wedge \ isFinal_{\mathsf{alg}}(res_1) \\ res_2 & \texttt{elseif } \otimes\mathsf{alg}(res_1, \mathcal{P}[\![p_2]\!]r) = res_2 \\ & \qquad \wedge \ isFinal_{\mathsf{alg}}(res_2) \\ \vdots & \qquad \vdots \\ res_{s\text{-}1} & \texttt{elseif } \otimes\mathsf{alg}(res_{s\text{-}2}, \mathcal{P}[\![p_{s\text{-}1}]\!]r) = res_{s\text{-}1} \\ & \qquad \wedge \ isFinal_{\mathsf{alg}}(res_{s\text{-}1}) \\ res_s & \texttt{otherwise } (\text{where } res_s = \otimes\mathsf{alg}(res_{s\text{-}1}, \mathcal{P}[\![p_s]\!]r) \end{cases} \end{aligned} \tag{S-6b}$$

[6] In case of a single policy, operators $\otimes$p-unless-d and $\otimes$d-unless-p turn the not-app and indet responses into, respectively, $\langle$permit $\epsilon\rangle$ and $\langle$deny $\epsilon\rangle$, while the remaining operators leave them unchanged.

where the `elseif` notation is a shortcut to represent mutually exclusive conditions. The auxiliary predicates $isFinal_{\mathsf{alg}}$ (one for each combining algorithm alg), given a response in input, check if the response decision is final with respect to the algorithm alg, i.e. if such decision cannot change due to further combinations. Their definition is in Table 6(b); as a matter of notation, we use $res.dec$ to indicate the decision of response res. These predicates are straightforwardly derived from the combination matrices of the binary operators, thus we only comment on salient points. In case of the p-over algorithm (and similarly for the others in the first two rows of the table), the permit decision is the only decision that can never be overwritten, hence, it is final. In case of the first-app algorithm, instead, all decisions except not-app are final since they represent the fact that the first applicable policy has been already found. Both consensus algorithms have indet as final decision, because no form of consensus can be reached once an indet is obtained. Similarly, the one-app algorithm has indet as final decision.

5.4 Semantics of the Policy Enforcement Process

The semantics of the enforcement process defines how the PEP discharges obligations and enforces authorisation decisions. To define this process, we use the auxiliary function $(\!(\ )\!) : IOblgation^* \to \{\mathsf{true}, \mathsf{false}\}$ that, given a sequence of instantiated obligations, executes such obligations and returns a boolean value that indicates whether the evaluation is successfully completed. Since failures caused by optional obligations can be safely ignored by the PEP, only failures of mandatory obligations (i.e. of type m) have to be taken into account. The function is defined as follows

$$\begin{aligned}
(\!(\epsilon)\!) &= \mathsf{true} \\
(\!([\,\mathsf{o}\ \ pepAct(w^*)\,] \bullet io^*)\!) &= (\!(io^*)\!) \\
(\!([\,\mathsf{m}\ \ pepAct(w^*)\,] \bullet io^*)\!) &= \begin{cases} (\!(io^*)\!) & \texttt{if } pepAct(w^*) \Downarrow \mathsf{ok} \\ \mathsf{false} & \texttt{otherwise} \end{cases}
\end{aligned}$$

where $\Downarrow\mathsf{ok}$ denotes if the discharge of the action $pepAct(w^*)$ succeeded. Since the set of action identifiers is intentionally left unspecified (see Section 4.1), the definition of the predicate $\Downarrow\mathsf{ok}$ is hence unspecified too; we just assume that it is total and deterministic. In other words, the syntactic domain *PepAction* is a parameter of the syntax, while the predicate $\Downarrow\mathsf{ok}$ is a parameter of the semantics. The latter parameter could be refined to deal with, e.g., obligations to be enforced after the decision releasing (see Section 10). For example, discharging obligations could simply refer to the fact that the system has taken charge of their execution, rather than to the fact that they have been completely executed.

The semantics of PEP is thus defined with respect to the enforcement algorithms. Formally, given an enforcement algorithm and a PDP response, the function $\mathcal{E}A : EnfAlg \to (PDPReponse \to Decision)$ returns the enforced decision. It is defined by three clauses, one for each algorithm. The clause for the deny-biased algorithm follows

$$\mathcal{E}A[\![\mathsf{deny\text{-}biased}]\!]res = \begin{cases} \mathsf{permit} & \texttt{if } res.dec = \mathsf{permit} \ \wedge\ (\!(res.io)\!) \\ \mathsf{deny} & \texttt{otherwise} \end{cases} \tag{S-7a}$$

Likewise $res.dec$ that indicates the decision of the response res, notation $res.io$ indicates the sequence of instantiated obligations of res. The permit decision is enforced only if this is the decision returned by the PDP and all accompanying obligations are successfully discharged. If an error occurs, as well as if the PDP decision is not permit, a deny is enforced. The clause for the permit-biased algorithm is the dual one, whereas the clause for the base algorithm is as follows

$$\mathcal{E}A[\![\mathsf{base}]\!]res = \begin{cases} \mathsf{permit} & \texttt{if } res.dec = \mathsf{permit} \ \wedge\ (\!(res.io)\!) \\ \mathsf{deny} & \texttt{if } res.dec = \mathsf{deny} \ \wedge\ (\!(res.io)\!) \\ \mathsf{not\text{-}app} & \texttt{if } res.dec = \mathsf{not\text{-}app} \\ \mathsf{indet} & \texttt{otherwise} \end{cases} \tag{S-7b}$$

Both decisions permit and deny are enforced only if all obligations in the PDP response are successfully discharged, otherwise they are enforced as indet. Instead, decisions not-app and indet are enforced without modifications.

5.5 Semantics of the Policy Authorisation System

The semantics of a Policy Authorisation System is defined in terms of the composition of the semantics of PEP and PDP. It is given by the function $\mathcal{P}as : PAS \to (Request \to Decision)$ defined by the following clause

$$\mathcal{P}as[\![\{\ \mathsf{pep} : ea \quad \mathsf{pdp} : pdp\ \}, req]\!] = \mathcal{E}A[\![ea]\!](\mathcal{P}dp[\![pdp]\!](\mathcal{R}[\![req]\!])) \tag{S-8}$$

Basically, given a request *req* in the FACPL syntax, this is converted into its functional representation by the function $\mathcal{R}$ (see Section 5.1). This result is then passed to the semantics of the PDP, i.e. $\mathcal{P}dp[\![pdp]\!]$, which returns a response that on its turn is passed to the semantics of the PEP, i.e. $\mathcal{E}A[\![ea]\!]$. The latter function returns the final decision of the Policy Authorisation System when given the request *req* in input.

5.6 Properties of the Semantics

We conclude this section with some properties and results regarding the FACPL semantics.

The main result is that the semantics is *total* and *deterministic*. This means that it is defined for all possible input pairs consisting of a FACPL specification, i.e. a Policy Authorisation System, and a request, and that it always returns the same decision any time it is applied to a specific pair.

Theorem 5.1 (Total and Deterministic Semantics)**.**

1. *For all $pas \in PAS$ and $req \in Request$, there exists a $dec \in Decision$, such that $\mathcal{P}as[\![pas, req]\!] = dec$.*

2. *For all $pas \in PAS$, $req \in Request$ and $dec, dec' \in Decision$, it holds that*
   $\mathcal{P}as[\![pas, req]\!] = dec \quad \wedge \quad \mathcal{P}as[\![pas, req]\!] = dec'$
   $\Rightarrow \quad dec = dec'$.

Proof. It boils down to show that $\mathcal{P}as$ is a total and deterministic function (see Appendix B.1). □

We now consider the so-called *reasonability properties* of [45] that precisely characterise the expressiveness of a policy language. FACPL enjoys the property called *independent composition* of policies, which means that the results of the combining algorithms depend only on the decisions of the policies given in input. This clearly follows from the use of combination matrices. On the contrary, FACPL ensures neither *safety*, i.e. a request that is granted may not be granted anymore if it is extended with new attributes, nor *monotonicity*, i.e. the introduction of a new policy in a combination of policies may change a permit decision to a different one. This should be somehow expected as these latter two properties are enjoyed neither by XACML nor by other policy languages featuring deny rules and combining algorithms similar to those we have presented.

We conclude by highlighting the relationship between attribute names occurring in a policy and names defined by requests. By letting $Names(p)$ to indicate the set of attribute names occurring in (the expressions within) p, we can state the following result which has important practical implications on the feasibility of the automatic analysis.

Lemma 5.2 (Policy relevant attributes)**.** *For all $p \in Policy$ and $r, r' \in R$ such that $r(n) = r'(n)$ for all $n \in Names(p)$, it holds that $\mathcal{P}[\![p]\!]r = \mathcal{P}[\![p]\!]r'$.*

Proof. The property straightforwardly derives from the semantics of FACPL expressions and from Theorem 5.1 (see Appendix B.1). □

6 FACPL Constraint-based Representation

The analysis of access control policies is essential for ensuring confidentiality and integrity of system resources. In the case of FACPL, the analysis is made difficult by the hierarchical structure of policies, the presence of conflict resolution strategies and the intricacies deriving from the many involved controls. Moreover, no off-the-shelf analysis tool directly takes FACPL specifications in input. Hence, for enabling the analysis of FACPL policies through well-established and efficient software tools, we introduce a constraint formalism that permits, on the one hand, to uniformly represent policies and, on the other hand, to perform extensive checks of (a possibly infinite number of) requests.

The constraint-based representation we propose specifies satisfaction problems in terms of formulae based on multiple theories as, e.g., boolean and linear arithmetics. Such kind of formulae are usually called *satisfiability modulo theories* (SMT) formulae. The SMT-based approach is supported by the relevant progress made in the development of automatic SMT solvers (e.g., Z3 [12], CVC4 [4], Yices [15]), which make SMT formulae to be extensively employed in diverse analysis applications [13].

This section introduces our constraint-based representation of FACPL policies, while the analysis it enables is presented in Section 7. We first introduce the constraint formalism (Section 6.1), then we present the constraint representation of FACPL policies (Section 6.2) and some crucial results stating that it is a semantic-preserving representation (Section 6.3), and finally we show some examples of constraints obtained from our e-Health case study (Section 6.4).

6.1 A Constraint Formalism

The constraint formalism we present here extends boolean and inequality constraints with a few additional operators aiming at precisely representing FACPL constructs. Intuitively, a constraint is a relation defined through some conditions on a set of attribute names[7]. An assignment of values to attribute names satisfies a constraint if all constraint conditions are matched. Our formalism, besides usual operators and values, explicitly considers the role of missing attributes, by assigning $\perp$ to attribute names, and of run-time errors, i.e. type mismatches in constraint evaluations. In fact, according to the usually accepted semantics of access control policies (besides XACML, see, e.g., [8, 10]), a condition involving a missing attribute should not be evaluated to false by default.

Syntax. Constraints are written according to the following grammar.

$$
\begin{aligned}
\mathit{Constr} ::= {} & \mathit{Value} \mid \mathit{Name} \mid \mathtt{isMiss}(\mathit{Constr}) \\
& \mid \mathtt{isErr}(\mathit{Constr}) \mid \mathtt{isBool}(\mathit{Constr}) \\
& \mid \neg\, \mathit{Constr} \mid \dot{\neg}\, \mathit{Constr} \mid \mathit{Constr}\ \mathsf{cop}\ \mathit{Constr} \\
\mathsf{cop} ::= {} & \wedge \mid \vee \mid \dot{\wedge} \mid \dot{\vee} \mid = \mid > \mid \in \\
& \mid + \mid - \mid * \mid /
\end{aligned}
$$

where the nonterminals *Value* and *Name* are defined in Table 2. Thus, a constraint can be a literal value, an attribute name, or a more complex constraint obtained through predicates `isMiss`(), `isErr`() and `isBool`(), or through boolean, comparison and arithmetic operators. The operators $\neg$, $\wedge$ and $\vee$ are the usual boolean ones, while $\dot{\neg}$, $\dot{\wedge}$ and $\dot{\vee}$ correspond to the 4-valued ones of FACPL expressions which implement the special management of $\perp$ and error values.

In the sequel, in addition to the notations of Table 4, we use the letter c to denote a generic element of the set of all constraints identified by the nonterminal *Constr*.

Semantics. The semantics of constraints is modelled by the function $\mathcal{C} : \mathit{Constr} \to (R \to \mathit{Value} \cup 2^{\mathit{Value}} \cup \{\mathsf{error}, \perp\})$ inductively defined by the clauses in Table 7 (the clauses for $>$, $\in$, $-$, $*$ and $/$ are omitted as they are similar to those for $=$ or $+$). Hence, the semantics of a constraint is a function that, given the functional

[7] In the literature, constraints are typically defined on a set of *variables*. In our framework, the role of variables is played by attribute names. Therefore, to maintain a coherent terminology throughout the paper, we refer to constraint variables as attribute names.

Tab. 7: Semantics of constraints (T stands for one of the sets of literal values or for the powerset of the set of all literal values, and $i, j \in \{1, 2\}$ with $i \neq j$)

$$\mathcal{C}[\![n]\!]r = r(n) \qquad \mathcal{C}[\![v]\!]r = v$$

$$\mathcal{C}[\![\texttt{isMiss}(c)]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c]\!]r = \bot \\ \mathsf{false} & \texttt{otherwise} \end{cases} \qquad \mathcal{C}[\![\texttt{isErr}(c)]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c]\!]r = \mathsf{error} \\ \mathsf{false} & \texttt{otherwise} \end{cases} \qquad \mathcal{C}[\![\texttt{isBool}(c)]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c]\!]r \in \{\mathsf{true}, \mathsf{false}\} \\ \mathsf{false} & \texttt{otherwise} \end{cases}$$

$$\mathcal{C}[\![\dot{\neg}\ c]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c]\!]r = \mathsf{false} \\ \mathsf{false} & \texttt{if } \mathcal{C}[\![c]\!]r = \mathsf{true} \\ \bot & \texttt{if } \mathcal{C}[\![c]\!]r = \bot \\ \mathsf{error} & \texttt{otherwise} \end{cases} \qquad \mathcal{C}[\![c_1\ \dot{\wedge}\ c_2]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c_1]\!]r = \mathcal{C}[\![c_2]\!]r = \mathsf{true} \\ \mathsf{false} & \texttt{if } \mathcal{C}[\![c_1]\!]r = \mathsf{false} \texttt{ or } \mathcal{C}[\![c_2]\!]r = \mathsf{false} \\ \bot & \texttt{if } \mathcal{C}[\![c_i]\!]r = \bot \texttt{ and } \mathcal{C}[\![c_j]\!]r \in \{\mathsf{true}, \bot\} \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

$$\mathcal{C}[\![c_1\ \dot{\vee}\ c_2]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c_1]\!]r = \mathsf{true} \texttt{ or } \mathcal{C}[\![c_2]\!]r = \mathsf{true} \\ \mathsf{false} & \texttt{if } \mathcal{C}[\![c_1]\!]r = \mathcal{C}[\![c_2]\!]r = \mathsf{false} \\ \bot & \texttt{if } \mathcal{C}[\![c_i]\!]r = \bot \texttt{ and } \mathcal{C}[\![c_j]\!]r \in \{\mathsf{false}, \bot\} \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

$$\mathcal{C}[\![\neg\ c]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c]\!]r = \mathsf{false} \\ & \texttt{or } \mathcal{C}[\![c]\!]r = \bot \\ \mathsf{false} & \texttt{otherwise} \end{cases} \qquad \mathcal{C}[\![c_1\ \wedge\ c_2]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c_1]\!]r = \mathsf{true} \texttt{ and } \mathcal{C}[\![c_2]\!]r = \mathsf{true} \\ \mathsf{false} & \texttt{otherwise} \end{cases} \qquad \mathcal{C}[\![c_1\ \vee\ c_2]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c_1]\!]r = \mathsf{true} \texttt{ or } \mathcal{C}[\![c_2]\!]r = \mathsf{true} \\ \mathsf{false} & \texttt{otherwise} \end{cases}$$

$$\mathcal{C}[\![c_1\ =\ c_2]\!]r = \begin{cases} \mathsf{true} & \texttt{if } \mathcal{C}[\![c_1]\!]r, \mathcal{C}[\![c_2]\!]r \in T \texttt{ and } \mathcal{C}[\![c_1]\!]r = \mathcal{C}[\![c_2]\!]r \\ \mathsf{false} & \texttt{if } \mathcal{C}[\![c_1]\!]r, \mathcal{C}[\![c_2]\!]r \in T \texttt{ and } \mathcal{C}[\![c_1]\!]r \neq \mathcal{C}[\![c_2]\!]r \\ \bot & \texttt{if } \mathcal{C}[\![c_i]\!]r = \bot \texttt{ and } \mathcal{C}[\![c_j]\!]r \neq \mathsf{error} \\ \mathsf{error} & \texttt{otherwise} \end{cases} \qquad \mathcal{C}[\![c_1\ +\ c_2]\!]r = \begin{cases} \mathcal{C}[\![c_1]\!]r + \mathcal{C}[\![c_2]\!]r & \texttt{if } \mathcal{C}[\![c_1]\!]r, \mathcal{C}[\![c_2]\!]r \in \mathit{Double} \\ \bot & \texttt{if } \mathcal{C}[\![c_i]\!]r = \bot \texttt{ and } \mathcal{C}[\![c_j]\!]r \neq \mathsf{error} \\ \mathsf{error} & \texttt{otherwise} \end{cases}$$

representation of a request (i.e., an assignment of values to attribute names), returns a literal value or a set of literal values or one of the special values $\bot$ and error.

The semantics of constraints, except for the cases of predicates and usual boolean operators, mimics the semantic definitions of the corresponding FACPL expression operators defined in Table 5 (e.g., the constraint operator $\dot{\vee}$ corresponds to the expression operator or, as well as $+$ corresponds to add). The clause defining the semantics of predicate $\texttt{isMiss}(c)$ (resp. $\texttt{isErr}(c)$) returns true only if the constraint c evaluates to $\bot$ (resp. error), while that of predicate $\texttt{isBool}(c)$ returns true only if the constraint c evaluates to a boolean value. The clauses for usual boolean operators are instead defined by ensuring that only boolean values can be returned. Specifically, they explicitly define conditions leading to result true, while in all the other cases the result is false. The constraint $\neg\, c$ evaluates to true not only when the evaluation of c returns false, but also when it returns $\bot$. This is particularly convenient for translating FACPL policies because, in case of not-app decisions, $\bot$ is treated as false.

6.2 From FACPL Policies to Constraints

The constraint-based representation of a FACPL policy is a logical combination of the constraints representing targets, obligations and combining algorithms occurring within the policy. Of course, combining algorithms using the greedy instantiation strategy are not dealt with, as we cannot statically predict when the (sequential) evaluation of a sequence of policies can stop since the decision, that would have resulted from evaluating the whole sequence, has been obtained. The translation is formally, and *compositionally*, defined by a family of translation functions $\mathcal{T}.$, that return the constraints representing the different FACPL terms. We use the emphatic brackets $\{\!|$ and $|\!\}$ to represent the application of a translation function to a syntactic term.

We start by presenting the translation of FACPL expressions, whose operators are very close to (some of) those on constraints. The translation is formally given by the function $\mathcal{T}_E : \mathit{Expr} \to \mathit{Constr}$, whose defining

clauses are given below

$$\begin{aligned}
&\mathcal{T}_E\{\!|v|\!\} = v && \mathcal{T}_E\{\!|n|\!\} = n\\
&\mathcal{T}_E\{\!|\mathsf{not}(expr)|\!\} = \dot{\neg}\mathcal{T}_E\{\!|expr|\!\} && \text{(T-1)}\\
&\mathcal{T}_E\{\!|\mathsf{op}(expr_1, expr_2)|\!\} = \mathcal{T}_E\{\!|expr_1|\!\}\ \mathtt{getCop}(\mathsf{op})\ \mathcal{T}_E\{\!|expr_2|\!\}
\end{aligned}$$

Thus, $\mathcal{T}_E$ acts as the identity function on attribute names and values, and as an homomorphism on operators. In fact, FACPL negation corresponds to the constraint operator $\dot{\neg}$, while the binary FACPL operators correspond to the constraint operators returned by the auxiliary function `getCop`(). Its definition is straightforward, the main cases are defined as follows

$$\begin{aligned}
&\mathtt{getCop}(\mathsf{and}) = \dot{\wedge} && \mathtt{getCop}(\mathsf{or}) = \dot{\vee}\\
&\mathtt{getCop}(\mathsf{equal}) = = && \mathtt{getCop}(\mathsf{in}) = \in\\
&\mathtt{getCop}(\mathsf{greater\text{-}than}) = > && \mathtt{getCop}(\mathsf{add}) = +
\end{aligned}$$

The translation of (sequences of) obligations returns a constraint whose satisfiability corresponds to the successful instantiation of all the input obligations. The translation function $\mathcal{T}_{Ob} : \mathit{Obligation}^* \to \mathit{Constr}$ is defined below

$$\begin{aligned}
&\mathcal{T}_{Ob}\{\!|\epsilon|\!\} = \mathsf{true}\\
&\mathcal{T}_{Ob}\{\!|o\ o^*|\!\} = \mathcal{T}_{Ob}\{\!|o|\!\} \wedge \mathcal{T}_{Ob}\{\!|o^*|\!\} && \text{(T-2)}\\
&\mathcal{T}_{Ob}\{\!|[t\ \mathit{PepAction}(expr^*)]|\!\} = \textstyle\bigwedge_{expr \in expr^*} \neg\mathtt{isMiss}(\mathcal{T}_E\{\!|expr|\!\}) \wedge \neg\mathtt{isErr}(\mathcal{T}_E\{\!|expr|\!\})
\end{aligned}$$

Hence, a sequence of obligations corresponds to the conjunction of the constraints representing each obligation. When translating a single obligation, predicates `isMiss`() and `isErr`() are used to check the instantiation conditions, i.e. that the occurring expressions cannot evaluate to $\bot$ or error. The n-ary conjunction operator returns true if the considered obligation contains no expression (i.e., $expr^* = \epsilon$).

The translation function for policies, $\mathcal{T}_P$, exploits the translation functions previously introduced, as well as a function $\mathcal{T}_A$ representing the result of applying a combining algorithm to a sequence of policies. Functions $\mathcal{T}_P$ and $\mathcal{T}_A$ are indeed mutually recursive. Moreover, for representing all the decisions that a policy can return, both these two functions return 4-tuples of constraints of the form

$$\langle \mathsf{permit} : c_p \quad \mathsf{deny} : c_d \quad \mathsf{not\text{-}app} : c_n \quad \mathsf{indet} : c_i \ \rangle$$

where each constraint represents the conditions under which the corresponding decision is returned. We call these tuples *policy constraint tuples* and denote their set by PCT. As a matter of notation, we will use the projection operator $\downarrow_l$ which, when applied to a constraint tuple, returns the value of the field labelled by l', where l is the first letter of l' (e.g., $\downarrow_p$ returns the permit constraint c_p).

The function $\mathcal{T}_P : \mathit{Policy} \to \mathit{PCT}$ is defined by two clauses for rules, i.e. one for each effect, and one clause for policy sets. The clause for rules with effect permit is

$$\begin{aligned}
&\mathcal{T}_P\{\!|(\mathsf{permit}\ \ \mathsf{target} : expr\ \ \mathsf{obl} : o^*\ )|\!\} =\\
&\ \langle\ \mathsf{permit} :\ \mathcal{T}_E\{\!|expr|\!\} \wedge\ \mathcal{T}_{Ob}\{\!|o^*|\!\}\\
&\ \ \mathsf{deny} :\ \mathsf{false} && \text{(T-3a)}\\
&\ \ \mathsf{not\text{-}app} :\ \neg\ \mathcal{T}_E\{\!|expr|\!\}\\
&\ \ \mathsf{indet} : \neg\,(\mathtt{isBool}(\mathcal{T}_E\{\!|expr|\!\})\ \vee\ \mathtt{isMiss}(\mathcal{T}_E\{\!|expr|\!\})) \vee\ (\mathcal{T}_E\{\!|expr|\!\} \wedge \neg\,\mathcal{T}_{Ob}\{\!|o^*|\!\})\ \rangle
\end{aligned}$$

(the clause for effect deny is omitted, as it only differs from the previous one because it swaps the permit and deny constraints). The clause takes into account the rule constituent parts and combines them according to the rule semantics (see clause (S-4a)). Because of the semantics of the constraint operator $\neg$, the not-app constraint is satisfied when the constraint corresponding to the target expression evaluates to false or to $\bot$. Instead, the negation of a constraint corresponding to a sequence of obligations represents the failure

of their instantiation. In the indet constraint, together with condition $\neg$ $\texttt{isBool}(\mathcal{T}_E\{\!|expr|\!\})$, we introduce $\neg$ $\texttt{isMiss}(\mathcal{T}_E\{\!|expr|\!\})$ because we want to exclude that $\mathcal{T}_E\{\!|expr|\!\} = \bot$ (otherwise, we would fall in the case of decision not-app).

The clause for policy sets is as follows

$$
\begin{array}{l}
\mathcal{T}_P\{\!|\langle\ a\ \ \mathsf{target}: expr\ \ \mathsf{policies}: p^+\ \mathsf{obl\text{-}p}: o_p^*\ \ \mathsf{obl\text{-}d}: o_d^*\ \rangle|\!\} \\
= \\
\langle\ \mathsf{permit}:\ \mathcal{T}_E\{\!|expr|\!\} \wedge \mathcal{T}_A\{\!|a, p^+|\!\} \downarrow_p \wedge \mathcal{T}_{Ob}\{\!|o_p^*|\!\} \\
\quad \mathsf{deny}:\ \mathcal{T}_E\{\!|expr|\!\} \wedge \mathcal{T}_A\{\!|a, p^+|\!\} \downarrow_d \wedge \mathcal{T}_{Ob}\{\!|o_d^*|\!\} \\
\quad \mathsf{not\text{-}app}:\ \neg\ \mathcal{T}_E\{\!|expr|\!\} \vee (\mathcal{T}_E\{\!|expr|\!\} \wedge \mathcal{T}_A\{\!|a, p^+|\!\} \downarrow_n) \\
\quad \mathsf{indet}: \\
\qquad \neg\,(\texttt{isBool}(\mathcal{T}_E\{\!|expr|\!\}) \vee \texttt{isMiss}(\mathcal{T}_E\{\!|expr|\!\})) \\
\qquad \vee\ (\mathcal{T}_E\{\!|expr|\!\} \wedge \mathcal{T}_A\{\!|a, p^+|\!\} \downarrow_i) \\
\qquad \vee\ (\mathcal{T}_E\{\!|expr|\!\} \wedge \mathcal{T}_A\{\!|a, p^+|\!\} \downarrow_p \wedge \neg\ \mathcal{T}_{Ob}\{\!|o_p^*|\!\}) \\
\qquad \vee\ (\mathcal{T}_E\{\!|expr|\!\} \wedge \mathcal{T}_A\{\!|a, p^+|\!\} \downarrow_d \wedge \neg\ \mathcal{T}_{Ob}\{\!|o_d^*|\!\})\ \ \rangle
\end{array}
\tag{T-3b}
$$

With respect to the clauses for rules, it additionally takes into account the result of the application of the combining algorithm according to the policy set semantics (see clause (S-4b)). It is worth noticing that the exclusive use of operators $\neg$, $\wedge$ and $\vee$ ensures that constraint tuples are only formed by boolean constraints.

Combining algorithms are dealt with by the function $\mathcal{T}_A : Alg \times Policy^+ \to PCT$ that, given an algorithm (using the all instantiation strategy) and a sequence of policies, returns a constraint tuple representing the result of the algorithm application. Its definition is

$$
\begin{array}{l}
\mathcal{T}_A\{\!|\mathsf{alg}_{\mathsf{all}}, p_1\ \ldots\ p_s|\!\} = \\
\quad \mathsf{alg}(\ldots \mathsf{alg}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\}), \ldots, \mathcal{T}_P\{\!|p_s|\!\})
\end{array}
\tag{T-4}
$$

By means of $\mathcal{T}_P$, the policies given in input are translated into constraint tuples which are then iteratively combined, two at a time, according to the algorithm combination strategy. By way of example, the combination of two constraint tuples, say A and B, according to the p-over algorithm, is defined as follows

$$
\begin{array}{l}
\mathsf{p\text{-}over}(A, B) = \\
\quad \langle \mathsf{permit}: \quad A \downarrow_p \vee B \downarrow_p \\
\quad\ \mathsf{deny}: \quad (A \downarrow_d \wedge B \downarrow_d) \vee (A \downarrow_d \wedge B \downarrow_n) \vee (A \downarrow_n \wedge B \downarrow_d) \\
\quad\ \mathsf{not\text{-}app}: \quad A \downarrow_n \wedge B \downarrow_n \\
\quad\ \mathsf{indet}: \quad (A \downarrow_i \wedge \neg B \downarrow_p) \vee (\neg A \downarrow_p \wedge B \downarrow_i) \rangle
\end{array}
$$

The combinations for the remaining algorithms are in Appendix A. If $s = 1$, i.e. there is only one argument tuple, all the algorithms leave the input tuple unchanged, but for p-unless-d, which given an input tuple A returns the tuple

$$
\begin{array}{ll}
\langle\ \mathsf{permit}: A \downarrow_p \vee A \downarrow_n \vee A \downarrow_i & \mathsf{deny}: A \downarrow_d \\
\quad \mathsf{not\text{-}app}: \mathsf{false} & \mathsf{indet}: \mathsf{false}\ \rangle
\end{array}
$$

and d-unless-p, which behaves similarly.

Finally, the translation of top-level PDP terms $\{Alg\ \ \mathsf{policies}: Policy^+\}$ is the same as that of the corresponding policy sets with target true and no obligations, i.e. $\{Alg\ \ \mathsf{target}: \mathsf{true}\ \ \mathsf{policies}: Policy^+\}$.

6.3 Properties of the Translation

The key result regarding the translation is that the semantics of the constraint-based representation of a policy and the semantics of the policy itself do agree. This correspondence is clearly limited to only those policies using the instantiation strategy all. Before presenting this result, we show for the constraint semantics a result analogous to Theorem 5.1.

Theorem 6.1 (Total and Deterministic Constraint Semantics)**.**

1. *For all $c \in Constr$ and $r \in R$, there exists an $el \in (Value \cup 2^{Value} \cup \{\mathsf{error}, \bot\})$, such that $\mathcal{C}[\![c]\!]r = el$.*

2. *For all $c \in Constr$, $r \in R$ and $el, el' \in (Value \cup 2^{Value} \cup \{\mathsf{error}, \bot\})$, it holds that*

$$\mathcal{C}[\![c]\!]r = el \quad \wedge \quad \mathcal{C}[\![c]\!]r = el' \quad \Rightarrow \quad el = el' .$$

Proof. By structural induction on the syntax of c (see Appendix B.2). □

Theorem 6.2 (Policy Semantic Correspondence)**.** *For all $p \in Policy$ enclosing combining algorithms only using* all *as instantiation strategy, and $r \in R$, it holds that*

$$\mathcal{P}[\![p]\!]r = \langle dec\ io^* \rangle \quad \Leftrightarrow \quad \mathcal{C}[\![\mathcal{T}_P\{\!|p|\!\}\downarrow_{dec}]\!]r = \mathsf{true}$$

Proof. The proof (see Appendix B.2) is by induction on the *depth*, i.e. the nesting level, of p and relies on three auxiliary correspondence results regarding expressions (Lemma B.1), obligations (Lemma B.2) and combining algorithms (Lemma B.3). □

This theorem implies that the properties verified over the constraints resulting from the translation of a FACPL policy would return the same results as if they were directly proven on the FACPL policy itself. Thus, it ensures that the analysis we present in Section 7 is sound.

From the previous theorems it follows that policy constraint tuples partition the set of input requests, in other words each access request satisfies only one of the constraints of a policy constraint tuple. Essentially, the following corollary extends Theorem 6.1 to constraint tuples.

Corollary 6.3 (Constraint-based partition)**.** *For all $r \in R$ and $p \in Policy$, such that $\mathcal{T}_P\{\!|p|\!\} = \langle$*permit *: c_1* deny *: c_2* not-app *: c_3* indet *: $c_4 \rangle$, it holds that*

$$\exists! k \in \{1, \ldots, 4\} \ : \mathcal{C}[\![c_k]\!]r = \mathsf{true} \ \wedge \ \bigwedge_{j \in \{1,\ldots,4\} \setminus \{k\}} \mathcal{C}[\![c_j]\!]r = \mathsf{false}$$

Proof. The thesis immediately follows from Theorems 6.1 and 6.2. □

6.4 Constraint-based Representation of the e-Health case study

We now apply the translation functions introduced in Section 6.2 to (a part of) the considered case study. For the sake of presentation, we shorten the attribute names used within policies. For instance, the rule addressing Requirement (1) becomes as follows

$$\begin{array}{ll} (\,\mathsf{permit} & \mathsf{target}: \mathsf{equal}(\mathsf{sub/role}, \text{“doctor”}) \\ & \mathsf{and}\ \mathsf{equal}(\mathsf{act/id}, \text{“write”}) \\ & \mathsf{and}\ \mathsf{in}(\text{“e-Pre-Write”}, \mathsf{sub/perm}) \\ & \mathsf{and}\ \mathsf{in}(\text{“e-Pre-Read”}, \mathsf{sub/perm}))\end{array}$$

Its translation starts by applying function $\mathcal{T}_E$ to the target expression. The resulting constraint is as follows

$$\begin{array}{rl} c_{trg1} \triangleq & \mathsf{sub/role} = \text{“doctor”} \ \dot{\wedge}\ \mathsf{act/id} = \text{“write”} \\ & \dot{\wedge}\ \text{“e-Pre-Write”} \in \mathsf{sub/perm} \\ & \dot{\wedge}\ \text{“e-Pre-Read”} \in \mathsf{sub/perm}\end{array}$$

The translation proceeds by considering obligations; in this case they are missing (i.e., they correspond to the empty sequence ϵ), hence the constraint true is obtained. Function $\mathcal{T}_P$ finally defines the constraint tuple for the rule as follows

$$\begin{array}{l} \langle \mathsf{permit}: \ c_{trg1} \wedge \mathsf{true} \\ \ \mathsf{deny}: \ \mathsf{false} \\ \ \mathsf{not\text{-}app}: \ \neg c_{trg1} \\ \ \mathsf{indet}: \neg(\mathtt{isBool}(c_{trg1}) \vee \mathtt{isMiss}(c_{trg1})) \vee (c_{trg1} \wedge \neg\mathsf{true}) \rangle \end{array}$$

The tuples for the rules addressing Requirements (2) and (3) are defined similarly, they only differ in the constraints representing their targets, which are denoted as c_{trg2} and c_{trg3}, respectively.

We can now define the constraint-based representation of Policy (P1). Besides the target expression, which is straightforwardly translated to the constraint $c_{trgP} \triangleq \mathsf{res/typ} = \text{“e-Pre”}$, the constraint tuple is built up from the result of function $\mathcal{T}_A$ representing the application of the algorithm p-over. Specifically, the constraint tuples of rules are iteratively combined according to the definition of p-over(A, B) previously reported. For example, the combination of the first two rules generates the following tuple

$$\begin{array}{l}\langle\, \mathsf{permit}: \ (c_{trg1} \wedge \mathsf{true}) \vee (c_{trg2} \wedge \mathsf{true}) \\ \mathsf{deny}: \ (\mathsf{false} \wedge \mathsf{false}) \vee (\mathsf{false} \wedge \neg c_{trg2}) \vee (\neg c_{trg_1} \wedge \mathsf{false}) \\ \mathsf{not\text{-}app}: \ \neg c_{trg1} \wedge \neg c_{trg2} \\ \mathsf{indet}: ((\neg(\mathtt{isBool}(c_{trg1}) \vee \mathtt{isMiss}(c_{trg1})) \\ \qquad \vee (c_{trg1} \wedge \neg\mathsf{true})) \wedge \neg(c_{trg2} \wedge \mathsf{true})) \\ \qquad \vee (\neg(c_{trg1} \wedge \mathsf{true}) \wedge \ (\neg(\mathtt{isBool}(c_{trg2}) \\ \qquad \vee \mathtt{isMiss}(c_{trg2})) \vee (c_{trg2} \wedge \neg\mathsf{true}))) \ \rangle\end{array}$$

Notably, the deny constraint is never satisfied, because it is a disjunction of conjunctions having at least one false term as argument. This is somewhat expected, because the rules have the permit effect and the used combining algorithm is p-over. This tuple is then combined with that of the remaining rule in a similar way.

To generate the constraint tuple of the policy, we also need the constraint-based representation of its obligations. The policy contains only one obligation for the effect permit, whose corresponding constraint is as follows

$$c_{obl_p} \triangleq \bigwedge_{n \in \{\mathsf{sys/time,res/typ,sub/id,act/id}\}} \neg\mathtt{isMiss}(n) \wedge \neg\mathtt{isErr}(n)$$

The constraint corresponding to obligations for the effect deny, which are missing, is instead true.

Finally, the constraint tuple of Policy (P1) generated by function $\mathcal{T}_P$ is as follows

$$\begin{array}{l}\langle \mathsf{permit}: c_{trgP} \wedge ((c_{trg1} \wedge \mathsf{true}) \vee (c_{trg2} \wedge \mathsf{true}) \vee (c_{trg3} \wedge \mathsf{true})) \wedge c_{obl_p} \\ \quad \mathsf{deny}: c_{trgP} \wedge ((((\mathsf{false} \wedge \mathsf{false}) \vee (\mathsf{false} \wedge \neg c_{trg2}) \vee (\neg c_{trg_1} \wedge \mathsf{false})) \wedge \mathsf{false}) \\ \qquad \vee (((\mathsf{false} \wedge \mathsf{false}) \vee (\mathsf{false} \wedge \neg c_{trg2}) \vee (\neg c_{trg_1} \wedge \mathsf{false})) \wedge \neg c_{trg3}) \\ \qquad \vee ((\neg c_{trg1} \wedge \neg c_{trg2}) \wedge \mathsf{false})) \wedge \ \mathsf{true} \\ \quad \mathsf{not\text{-}app}: \neg c_{trgP} \vee (c_{trgP} \wedge (\neg c_{trg1} \wedge \neg c_{trg2} \wedge \neg c_{trg3})) \\ \quad \mathsf{indet}: \neg(\mathtt{isBool}(c_{trgP}) \ \vee \ \mathtt{isMiss}(c_{trgP})) \\ \qquad \vee (c_{trgP} \wedge (((\neg(\mathtt{isBool}(c_{trg1}) \ \vee \ \mathtt{isMiss}(c_{trg1})) \vee (c_{trg1} \wedge \neg\mathsf{true})) \wedge \neg(c_{trg2} \wedge \mathsf{true})) \\ \qquad \vee \neg((c_{trg1} \wedge \mathsf{true}) \wedge (\neg(\mathtt{isBool}(c_{trg2}) \vee \mathtt{isMiss}(c_{trg2})) \\ \qquad\quad \vee (c_{trg2} \wedge \neg\mathsf{true}))) \wedge \neg(c_{trg3} \wedge \mathsf{true})) \vee \ (\neg((c_{trg1} \wedge \mathsf{true}) \vee (c_{trg2} \wedge \mathsf{true})) \wedge (\neg(\mathtt{isBool}(c_{trg3}) \\ \qquad\quad \vee \mathtt{isMiss}(c_{trg3})) \vee (c_{trg3} \wedge \neg\mathsf{true}))) \\ \qquad \vee (c_{trgP} \wedge ((c_{trg1} \wedge \mathsf{true}) \vee (c_{trg2} \wedge \mathsf{true}) \vee (c_{trg3} \wedge \mathsf{true})) \wedge \neg c_{obl_p}) \\ \qquad \vee (c_{trgP} \wedge ((((\mathsf{false} \wedge \mathsf{false}) \vee (\mathsf{false} \wedge \neg c_{trg2}) \vee (\neg c_{trg_1} \wedge \mathsf{false})) \wedge \mathsf{false}) \\ \qquad \vee (((\mathsf{false} \wedge \mathsf{false}) \vee (\mathsf{false} \wedge \neg c_{trg2}) \vee (\neg c_{trg_1} \wedge \mathsf{false})) \wedge \ \neg c_{trg3}) \\ \qquad \vee ((\neg c_{trg1} \wedge \neg c_{trg2}) \wedge \mathsf{false})) \wedge \neg\mathsf{true})\rangle\end{array}$$

As this example demonstrates, the constraints resulting from the translation are a single-layered representation of policies that fully details all the aspects of policy evaluation. However, it is also evident that the evaluation, as well as the generation, of such constraints cannot be done manually, but requires a tool support.

7 Analysis of FACPL Policies

The analysis of FACPL policies we propose aims at verifying different types of properties by exploiting the constraint-based representation of policies. We first formalise a relevant set of properties in terms of expected authorisations for requests, and then we define the strategies for their automated verification by means of constraints.

Furthermore, since FACPL does not enjoy the *safety* property (see Section 5.6), the analysis investigates how the extension of a request through the addition of further attributes might change its authorisation in a possibly unexpected way. Intuitively, it is important to consider the authorisation decisions not only of specific requests, but also of their extensions because, e.g., a malicious user could try to exploit them to circumvent the access control system. This analysis approach is partially inspired by the probabilistic analysis on missing attributes introduced in [9].

In the following, we first formalise the proposed properties (Section 7.1) and present some concrete examples of them from the case study (Section 7.2). Afterwards, we show how to express the constraint formalism into a tool-accepted specification (Section 7.3) and exploit it to automatically verify the properties with an SMT solver (Section 7.4).

7.1 Formalisation of Properties

We consider both properties that refer to the expected authorisation of single requests, i.e. *authorisation properties* (Section 7.1.1), and to the relationships among policies on the base of the whole set of authorisations they establish, i.e. *structural properties* (Section 7.1.2); afterwards we comment on their automatic verification (Section 7.1.3).

7.1.1 Authorisation Properties

To formalise the authorisations properties, we introduce the notion of *request extension set* of a given request r. It is defined as follows

$$Ext(r) \triangleq \{r' \in R \mid r(n) \neq \bot \quad \Rightarrow \quad r'(n) = r(n)\}$$

The set is formed by all those requests that possibly extend request r with new attributes not already defined by r.

Evaluate-To. This property, written r `eval` dec, requires the policy under examination to evaluate the request r to decision dec. The satisfiability, written `sat`, of the *Evaluate-To* property by a policy p is defined as follows

$$p \texttt{ sat } r \texttt{ eval } dec \qquad \textit{iff} \qquad \mathcal{P}[\![p]\!]r = \langle dec \quad io^* \rangle$$

In practice, the verification of the property boils down to apply the semantic function $\mathcal{P}$ to p and r, and to check that the resulting decision is dec.

May-Evaluate-To. This property, written r $\texttt{eval}_{\texttt{may}}$ dec, requires that *at least one* request extending the request r evaluates to decision dec. The satisfiability of the *May-Evaluate-To* property by a policy p is defined as follows

$$p \texttt{ sat } r \ \texttt{eval}_{\texttt{may}} \ dec \quad \textit{iff} \ \exists\, r' \in Ext(r) \ : \ \mathcal{P}[\![p]\!]r' = \langle dec \quad io^* \rangle$$

This property, as well as the next one, addresses additional attributes extending the request r by considering the requests in its extension set $Ext(r)$.

Must-Evaluate-To. This property, written r $\texttt{eval}_{\texttt{must}}$ dec, differs from the previous one as it requires *all* the extended requests to evaluate to decision dec. The satisfiability of the *Must-Evaluate-To* property by a policy p is defined as follows

$$p \texttt{ sat } r \ \texttt{eval}_{\texttt{must}} \ dec \qquad \textit{iff} \ \forall r' \in Ext(r) \ : \ \mathcal{P}[\![p]\!]r' = \langle dec \quad io^* \rangle$$

Of course, additional properties can be obtained by combining the previous ones like, e.g., a property requiring that all requests in $Ext(r)$ may evaluate to dec and must not evaluate to dec'. Again, request extensions can be exploited to track down possibly unexpected authorisations.

7.1.2 Structural Properties

A structural property refers to the structure of the sets of authorisations established by one or multiple policies. In case of multiple policies, the properties aim at characterising the relationships among the policies. Different structural properties have been proposed in the literature (e.g. in [18] and [26]) by pursuing different approaches for their definition and verification. Here, we consider a set of commonly addressed properties and provide a uniform characterisation thereof in terms of requests and policy semantics.

Completeness. A policy is `complete` if it applies to all requests. Thus, the satisfiability of the *Completeness* property by a policy p is defined as follows

$$p \texttt{ sat complete} \qquad \textit{iff } \forall\ r\ \in R\ :\ \mathcal{P}[\![p]\!]r = \langle dec\ \ io^* \rangle, dec \neq \mathsf{not\text{-}app}$$

Essentially, we require that the policy applies to any request, i.e. it always returns a decision different from not-app. Notably, in this formulation indet is considered as an acceptable decision; a more restrictive formulation could only accept permit and deny.

Disjointness. Disjointness among policies means that such policies apply to disjoint sets of requests. Thus, this property, written `disjoint` p', requires that there is no request for which both the policy under examination and the policy p' evaluate to a decision considered *admissible*, i.e. permit or deny. The satisfiability of the *Disjointness* property by a policy p is defined as follows

$$p \texttt{ sat disjoint } p' \qquad \textit{iff} \qquad \forall\ r \in R\ :\ \mathcal{P}[\![p]\!]r = \langle dec\ \ io^* \rangle, \mathcal{P}[\![p']\!]r = \langle dec'\ \ io'^* \rangle, \\ \{\ dec, dec'\ \} \not\subseteq \{\mathsf{permit}, \mathsf{deny}\}$$

It is worth noticing that disjoint polices can be combined with the assurance that the allowed or forbidden authorisations established by each of them are not in conflict, which simplifies the choice of the combining algorithm to be used.

Coverage. Coverage among policies means that one of such policies establishes the same decisions as the other ones. More specifically, the property `cover` p' requires that for each request r for which p' evaluates to an admissible decision, the policy under examination evaluates to the same decision. The satisfiability of the *Coverage* property by a policy p is defined as follows

$$p \texttt{ sat cover } p' \qquad \textit{iff} \qquad \forall\ r \in R\ :\ \mathcal{P}[\![p']\!]r = \langle dec\ \ io^* \rangle, dec \in \{\mathsf{permit}, \mathsf{deny}\} \\ \Rightarrow\ \ \mathcal{P}[\![p]\!]r = \langle dec\ \ io'^* \rangle$$

Thus, p calculates at least the same admissible decisions as p'. Consequently, if p' also covers p, the two policies establish exactly the same admissible authorisations.

These structural properties permit statically reasoning on the relationships among policies and support system designers in developing and maintaining policies. One technique they enable is the *change-impact analysis* [18]. This analysis examines policy modifications for discovering unintended consequences of such changes.

7.1.3 Towards Automated Verification

It is worth noticing that the analysis approach we propose is feasible in practice, although the involved sets of requests might be infinite, e.g. the request extension set of a given request and the set of all possible requests. Indeed, Lemma 5.2 implies that only the attribute names occurring within the policies of interest are relevant for their analysis, and these are finite in number; any other name cannot affect policy evaluation. For instance, to analyse a policy p, we must not consider the set R of all possible requests, but only the set of those requests whose domain is $Names(p)$, i.e. the finite set of attribute names occurring in p. This property paves the way for carrying out automated property verification by means of SMT solvers as described in Section 7.4.

7.2 Properties on the e-Health case study

By way of example, we address in terms of authorisation and structural properties the case of pharmacists willing to write an e-Prescription in the e-Health case study.

Given the patient consent policies in Section 4.3, i.e. Policies (P1) and (P2), we can verify whether they disallow the access to a pharmacist that wants to write an e-Prescription. To this aim, we define an *Evaluate-To* property[8] as follows

$$(\mathsf{sub/role}, \text{“pharmacist”})(\mathsf{act/id}, \text{“write”})(\mathsf{res/typ}, \text{“e-Pre”})\ \texttt{eval}\ \ \mathsf{deny} \tag{Pr1}$$

which requires that such request evaluates to deny. Alternatively, by exploiting request extensions, we can check if there exists a request for which a pharmacist acting on e-Prescription can be evaluated to not-app. This corresponds to the *May-Evaluate-To* property defined as follows

$$(\mathsf{sub/role}, \text{“pharmacist”})(\mathsf{res/typ}, \text{“e-Pre”})\ \ \texttt{eval}_{\texttt{may}}\ \ \mathsf{not\text{-}app} \tag{Pr2}$$

The verification of these properties with respect to Policy (P1) results in

$$\mathit{Policy}\ (P1)\ \ \texttt{unsat}\ \ (Pr1)\quad \mathit{Policy}\ (P1)\ \ \texttt{sat}\ \ (Pr2)$$

where unsat indicates that the policy does not satisfy the property. Indeed, as already pointed out in Section 4.3, each request assigning to act/id a value different from read evaluates to not-app, hence property (Pr1) is not satisfied while property (Pr2) holds. On the contrary, the verification with respect to Policy (P2) results in

$$\mathit{Policy}\ (P2)\ \ \texttt{sat}\ \ (Pr1)\quad \mathit{Policy}\ (P2)\ \ \texttt{unsat}\ \ (Pr2)$$

Both results are due to the internal policy (deny) which, together with the algorithm p-over, prevents not-app to be returned and establishes deny as default decision.

The analysis can also be conducted by relying on the structural properties. By verifying completeness, we can check if there is a request that evaluates to not-app. We get

$$\mathit{Policy}\ (P1)\ \ \texttt{unsat complete}\qquad \mathit{Policy}\ (P2)\ \ \texttt{sat complete}$$

As expected, Policy (P1) does not satisfy completeness, i.e. there is at least one request that evaluates to not-app, whereas Policy (P2) is complete. Instead, we can check if Policy (P2) correctly refines Policy (P1) by simply verifying coverage. We get

$$\mathit{Policy}\ (P2)\ \ \texttt{sat cover}\ \mathit{Policy}\ (P1)$$

This follows from the fact that Policy (P2) evaluates to permit the same set of requests as Policy (P1) and that Policy (P1) never returns deny; clearly, the opposite coverage property does not hold. It should be also noted that the two policies are not disjoint (as they share the set of permitted requests).

7.3 Expressing Constraints with SMT-LIB

Property verification requires extensive checks on large (possibly infinite) amounts of requests, hence, in order to be practically effective, tool support is essential. To this aim, we express the constraints defined in Section 6 by means of the SMT-LIB language (`http://smtlib.cs.uiowa.edu/`), that is a standardised constraint language accepted by most of the SMT solvers. Intuitively, SMT-LIB is a strongly typed functional language expressly defined for the specification of constraints. Of course, the feasibility of the SMT-based reasoning crucially depends on decidability of the satisfiability checks to be done; in other words, the used SMT-LIB constructs must refer to decidable theories, as e.g. uninterpreted function and array theories. We now provide a few insights on the SMT-LIB coding of our constraints.

The key element of the coding strategy is the parametrised record type representing attributes. This type, named `TValue`, is defined as follows

[8] For the sake of presentation, in this subsection we write requests using the FACPL syntax (i.e., they are specified as sequences of attributes) rather than using their semantical functional representation.

Tab. 8: Type inference rules for (an excerpt of) FACPL expressions; we use X as a type variable, U as a type name or a type variable, and we assume that *Bool*, *Double*, *String*, *Date*, 2^{Value} identify both the values' domains and their type names

$$\frac{v \in Bool}{\Gamma \vdash v : Bool \mid \mathsf{true}} \qquad \frac{v \in Double}{\Gamma \vdash v : Double \mid \mathsf{true}} \qquad \frac{v \in String}{\Gamma \vdash v : String \mid \mathsf{true}} \qquad \frac{v \in Date}{\Gamma \vdash v : Date \mid \mathsf{true}} \qquad \frac{v \in 2^{Value}}{\Gamma \vdash v : 2^{Value} \mid \mathsf{true}} \qquad \frac{\Gamma(n) = X}{\Gamma \vdash n : X \mid \mathsf{true}}$$

$$\frac{\Gamma \vdash expr : U \mid C}{\Gamma \vdash \mathsf{not}(expr) : Bool \mid C \wedge U = Bool} \qquad \frac{\Gamma \vdash expr_1 : U_1 \mid C_1 \qquad \Gamma \vdash expr_2 : U_2 \mid C_2}{\Gamma \vdash \mathsf{eop}(expr_1, expr_2) : Bool \mid C_1 \wedge C_2 \wedge U_1 = Bool \wedge U_2 = Bool} \; \mathsf{eop} \in \{\mathsf{and}, \mathsf{or}\}$$

$$\frac{\Gamma \vdash expr_1 : U_1 \mid C_1 \qquad \Gamma \vdash expr_2 : U_2 \mid C_2}{\Gamma \vdash \mathsf{equal}(expr_1, expr_2) : Bool \mid C_1 \wedge C_2 \wedge U_1 = U_2} \qquad \frac{\Gamma \vdash expr_1 : U_1 \mid C_1 \qquad \Gamma \vdash expr_2 : 2^{U_2} \mid C_2}{\Gamma \vdash \mathsf{in}(expr_1, expr_2) : Bool \mid C_1 \wedge C_2 \wedge U_1 = U_2}$$

```
(declare-datatypes (T) ((TValue
    (mk-val (val T)(miss Bool)(err Bool)))))
```

Hence, each attribute consists of a 3-valued record, whose first field `val` is the value with parametric type `T` assigned to the attribute, while the boolean fields `miss` and `err` indicate, respectively, if the attribute value is missing or has an unexpected type. Additional assertions, not shown here for the sake of presentation, ensure that the fields `miss` and `err` cannot be true at the same time, and that, when one of the last two fields is true, it takes precedence over `val`. Of course, a specification formed by multiple assertions is satisfied when all the assertions are satisfied.

The declaration of `TValue` outlines the syntax of SMT-LIB and its strongly typed nature. This means that each attribute occurring in a policy has to be typed, by properly instantiating the type parameter `T`. Since FACPL is an untyped language, to reconstruct the type of each attribute, we define the type inference system (whose excerpt is) reported in Table 8. The rules are straightforward and infer the judgment $\Gamma \vdash expr : U \mid C$ which, under the typing context Γ, assigns the type (or the type variable) U to the FACPL expression *expr* and generates the typing constraint C. Specifically, Γ is an injective function that associates a type variable to each attribute name, while C is basically made of conjunctions and disjunctions of equalities between variables and types. The generated typing constraint will be processed at the end of the inference process to establish well-typedness of an expression. Thus, a FACPL expression is *well-typed* if C is satisfiable, i.e. there exists a type assignment for the typing variables occurring in C that satisfies C. Moreover, a FACPL policy is *well-typed* if the typing constraints generated by all the expressions occurring in the policy are satisfied by a same assignment. These type assignments are then used to instantiate the type parameters of the SMT-LIB constraints representing well-typed policies.

The type inference system aims at statically getting rid of all those policies containing expressions that are not well-typed. For instance, given the expression $\mathsf{or}(\mathsf{cat/id}, \mathsf{equal}(\mathsf{cat/id}, 5))$ and the typing context $\Gamma(\mathsf{cat/id}) = X_{\mathsf{cat/id}}$, the inference rules assign the type *Bool* to the expression and generate the constraint $X_{\mathsf{cat/id}} = Double \wedge X_{\mathsf{cat/id}} = Bool \wedge Bool = Bool$. This constraint is clearly unsatisfiable (as attribute $\mathsf{cat/id}$ cannot simultaneously be a double and a boolean), hence a policy containing such expression is not well-typed and would be statically discarded. Notably, the use of the field `err` allows the analysis to address the role of errors in policy evaluations, i.e. to reason on the authorisations of requests assigning values of unexpected type to attribute names. It is indeed crucial to analyse also these requests, since a possible attacker can leverage on them to circumvent the access control system.

On top of the `TValue` datatype we build the uninterpreted functions expressing the operators of the proposed constraint formalism. By way of example, the operator $\dot{\wedge}$ corresponds to the `FAnd` function defined as follows

```
(define-fun FAnd
   ((x (TValue Bool)) (y (TValue Bool)))
   (TValue Bool)
   (ite (and (isTrue x) (isTrue y))
```

```
(mk-val true false false)
(ite (or (isFalse x) (isFalse y))
    (mk-val false false false)
    (ite (or (err x) (err y))
         (mk-val false false true)
         (mk-val false true  false)))))
```

where `mk-val` is the constructor of `TValue` records. Hence, the function takes as input two `TValue Bool` records (i.e., type `Bool` is the instantiation of the type parameter `T`) and returns a `Bool` record as well. The conditional if-then-else assertions `ite` are nested to form a structure that mimics the semantic conditions of Table 7, so that different `TValue` records are returned according to the input. The function `isFalse` (resp. `isTrue`) is used to compactly check that all fields of the record are false (resp. only the field `val` is true). All the other constraint operators, except $\in$, are defined similarly.

To express the operator $\in$, we need to represent multivalued attributes. Firstly, we define an array datatype, named `Set`, to model sets of elements as follows

```
(define-sort Set (T) (Array Int T))
```

where the type parameter `T` is the type of the elements of the array. By definition of array, each element has an associated integer index that is used to access the corresponding value. Thus, a multivalued attribute is represented by a `TValue` record with type an instantiated `Set`, e.g. `(TValue (Set Int))` is an attribute whose value is a set of integers. Consequently, we can build the uninterpreted function modelling the constraint operator $\in$. In case of integer sets, we have

```
(define-fun inInt
  ((x (TValue  Int)) (y (TValue (Set Int))))
  (TValue Bool)
  (ite (or (err x)(err y))
    (mk-val false false true)
    (ite (or (miss x) (miss y))
      (mk-val false true false)
      (ite (exists ((i Int))
             (=  (val x) (select (val y) i)))
           (mk-val true  false false)
           (mk-val false false false)))))
```

where the command `(select (val y) i)` takes the value in position `i` of the set in the field `val` of the argument `y`. In addition to the conditional assertions, the function uses the existential quantifier `exists` for checking if the value of the argument `x` is contained in the set of the argument `y`.

The coding approach we pursue generates, in most of the cases, fully decidable constraints. In fact, since we support non-linear arithmetic, i.e. multiplication, it is possible to define constraints for which a constraint solver is not able to answer. Anyway, modern constraint solvers are actually able to resolve nontrivial nonlinear problems that, for what concerns access control policies, should prevent any undefined evaluation[9]. Similarly, the quantifier-based constraints are in general not decidable, but solvers still succeed in evaluating complicated quantification assertions due to, e.g., powerful pattern techniques (see, e.g., the documentation of Z3). Notice anyway that if we assume that each expression operator in (and, consequently, constraint operator $\in$) is applied to at most one attribute name, the quantifications are bounded by the number of literals defining the other operator argument.

Concerning the value types we support, SMT-LIB does not provide a primitive type for *Date*. Hence, we use integers to represent its elements. Furthermore, even though SMT-LIB supports the *String* type, the Z3 solver we use does not. Thus, given a policy as an input, we define an additional datatype, say `Str`, with

[9] It should be noted that if at least one argument of each occurrence of the multiply operator is a numeric constant, then the resulting non-linear arithmetic constraints are decidable.

as many constants as the string values occurring in the policy. The string equality function is then defined over `TValue` records instantiated with type `Str`.

By way of example, the SMT-LIB code for the constraint c_{trg1} (see Section 6.4) is

```
(define-fun cns_target_Rule1
  ()
  (TValue Bool)
  (FAnd
    (equalStr n_sub/role cst_doc)
    (FAnd (equalStr n_act/id cst_write)
      (FAnd
        (inStr cst_permWrite n_sub/perm)
        (inStr cst_permRead n_sub/perm)))))
```

where identifiers starting with `n_` (resp. `cst_`) represent attribute names (resp. literals) of the represented expression. The whole SMT-LIB code for Policy (P1) can be found at `http://facpl.sf.net/eHealth/`.

7.4 Automated Properties Verification

The SMT-LIB coding permits using SMT solvers to automatically verify the properties formalised in Section 7.1. In the following, given a FACPL policy p, we denote by $\langle \mathsf{permit} : \mathit{smtlib}\text{-}c_p \ \mathsf{deny} : \mathit{smtlib}\text{-}c_d \ \mathsf{not\text{-}app} : \mathit{smtlib}\text{-}c_n \ \mathsf{indet} : \mathit{smtlib}\text{-}c_i \rangle$ the tuple of SMT-LIB codes representing the formal constraints $\mathcal{T}_P\{\!|p|\!\} = \langle \mathsf{permit} : c_p \ \mathsf{deny} : c_d \ \mathsf{not\text{-}app} : c_n \ \mathsf{indet} : c_i \rangle$. Hereafter, we present first the strategies to follow for verifying the authorisation properties, then those for verifying the structural properties.

7.4.1 Authorisation Properties

The verification of authorisation properties requires: (i) to introduce into the policy constraint of interest, which is chosen according to the property, the SMT-LIB coding of the request defined by the property; (ii) to check the satisfiability (or validity) of the resulting constraint.

Given a request r, the SMT-LIB coding of the request is defined as follows

$$r_{smtlib} \triangleq \left\{ \begin{array}{l} (\texttt{assert}\ (=\ (\texttt{val}\ \ n)\ \ v)) \\ (\texttt{assert}\ (\texttt{and}\ (\texttt{not}\ (\texttt{miss}\ \ n)) \\ \qquad\qquad (\texttt{not}\ (\texttt{err}\ \ n)))) \end{array} \middle|\ r(n) = v \right\}$$

Thus, all attribute names n in r are asserted to be equal to their value v and to be neither missing nor erroneous. Furthermore, given a FACPL policy p, we also define the following SMT-LIB coding of the request

$$\overline{r_{smtlib}(p)} \triangleq \left\{\ (\texttt{assert}\ (\texttt{miss}\ \ n))\ \middle|\ n \in \mathit{Names}(p)\ \wedge\ r(n) = \bot\ \right\}$$

where, as in Section 5.6, $\mathit{Names}(p)$ indicates the set of attribute names occurring in p. Thus, all the names n that occur in p and are not assigned to a value in r are asserted to be missing attributes.

By exploiting this SMT-LIB coding of requests, we define the automated verification (i.e., via an SMT solver) of the authorisation properties as follows

$$\begin{array}{l} p\ \texttt{sat}\ r\ \texttt{eval}\ \mathit{dec}\ \ \texttt{iff} \\ \qquad \mathit{smtlib}\text{-}c_{dec}\ \circ\ r_{smtlib}\ \circ\ \overline{r_{smtlib}(p)}\ \ \ \texttt{is sat} \\ p\ \texttt{sat}\ r\ \texttt{eval}_{\texttt{may}}\ \mathit{dec}\ \ \texttt{iff} \\ \qquad\qquad \mathit{smtlib}\text{-}c_{dec}\ \circ\ r_{smtlib}\ \ \texttt{is sat} \\ p\ \texttt{sat}\ r\ \texttt{eval}_{\texttt{must}}\ \mathit{dec}\ \texttt{iff} \\ \qquad\qquad \mathit{smtlib}\text{-}c_{dec}\ \circ\ r_{smtlib}\ \ \texttt{is valid} \end{array}$$

where $\circ$ indicates the concatenation of SMT-LIB code[10] and `valid` means that the corresponding SMT-LIB code is a valid set of assertions. Some comments follow.

The *Evaluate-To* property does not exploit request extensions, hence all attribute names not assigned by the considered request can only assume the special value $\perp$. This means that the request r is coded in SMT-LIB with r_{smtlib} and $\overline{r_{smtlib}(p)}$. The satisfiability of the property thus corresponds to that of the resulting SMT-LIB code.

To verify the *May-Evaluate-To* property, since it considers request extensions, the request has to be coded only with r_{smtlib}. As before, the satisfiability of the property corresponds to that of the resulting SMT-LIB code.

Finally, to verify the *Must-Evaluate-To* property, we code again the request with r_{smtlib} only, but we check the validity of the resulting SMT-LIB code, i.e. that it is satisfied by all the assignments for the attribute names. This amounts to check if the negation of the resulting SMT-LIB code is not satisfiable, in which case the property holds.

7.4.2 Structural Properties

The verification of structural properties does not require to modify policy constraints, but rather to check the unsatisfiability of combinations of constraints. It is defined as follows

$$p \texttt{ sat complete} \quad \texttt{iff} \quad \textit{smtlib-}c_n \quad \texttt{is unsat}$$

$$p \texttt{ sat disjoint } p' \texttt{ iff}$$
$$\begin{cases} \textit{smtlib-}c_p \circ \textit{smtlib-}c'_p & \texttt{is unsat} \\ \textit{smtlib-}c_p \circ \textit{smtlib-}c'_d & \texttt{is unsat} \\ \textit{smtlib-}c_d \circ \textit{smtlib-}c'_p & \texttt{is unsat} \\ \textit{smtlib-}c_d \circ \textit{smtlib-}c'_d & \texttt{is unsat} \end{cases}$$

$$p \texttt{ sat cover } p' \quad \texttt{iff}$$
$$\begin{cases} \neg\ \textit{smtlib-}c_p \circ \textit{smtlib-}c'_p & \texttt{is unsat} \\ \neg\ \textit{smtlib-}c_d \circ \textit{smtlib-}c'_d & \texttt{is unsat} \end{cases}$$

where $\textit{smtlib-}c'_{dec}$ refers to the SMT-LIB code modelling decision *dec* of policy p'. Some comments follow.

The trivial case is that of the *completeness* property, which only amounts to check if the constraint modelling the decision not-app is not satisfiable, i.e. if its negation is valid. If it is, the property holds.

The *disjointness* of two policies is verified by checking, one at a time, if the conjunctions between the permit or deny constraint of the first policy and the permit or deny constraint of the second policy are not satisfiable. If this holds for the four possible combinations of those constraints, the property holds.

The *coverage* of policy p on policy p' is verified by checking if the conjunction between the negation of the permit (resp., deny) constraint of p and the permit (resp., deny) constraint of p' is not satisfiable. Intuitively, if the policy p does not calculate a permit or deny decision (i.e., $\neg$ *smtlib-*c_p and $\neg$ *smtlib-*c_d hold), policy p' cannot do it as well, otherwise the property is not satisfied. If this holds for the two conjunctions separately, the property holds.

8 The FACPL Toolchain

The coding, analysis and enforcement tasks pursued in the development of FACPL specifications are fully supported by a Java-based software toolchain[11], graphically depicted in Figure 3. The key element of the toolchain is an Eclipse-based IDE that provides features like, e.g., static code checks and automatic generation of runnable Java and SMT-LIB code. An expressly developed Java library is used to compile and execute the Java code, while the analysis of SMT-LIB code exploits the Z3 solver.

[10] Notably, checking the satisfiability of the SMT-LIB code resulting from the concatenation of two (sets of) SMT-LIB assertions amounts to check if both the assertions hold at the same time.

[11] The FACPL supporting tools are freely available and open-source; binary files, source files, unit tests and documentation can be found at the FACPL website `http://facpl.sf.net`.

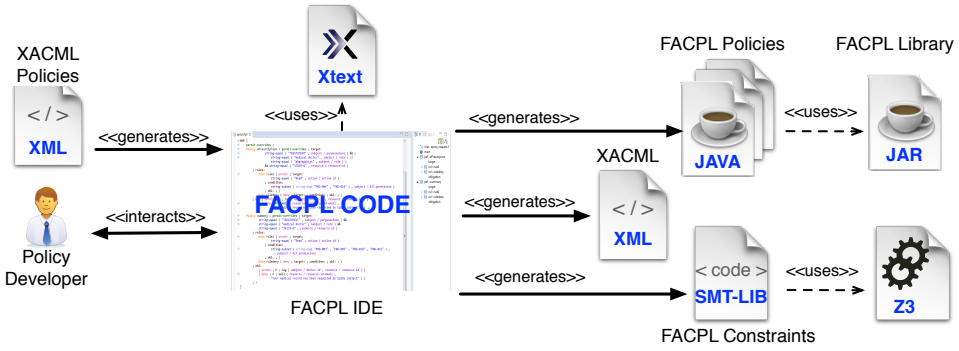

Fig. 3: The FACPL toolchain

To provide interoperability with the standard XACML and the various available tools supporting it (e.g., XCREATE [5], Margrave [18] and Balana [48]), the IDE automatically translates FACPL code into XACML one and vice-versa. Because of slightly different expressivity, there are some limitations in FACPL and XACML interoperability (see Section 9.1 for further details).

Furthermore, to allow newcomer users to directly experiment with FACPL, the web application "Try FACPL in your Browser" (reachable from the FACPL website) offers an online editor for creating and evaluating FACPL policies; the e-Health case study is there reported as a running example. Additionally, the web interface reachable from `http://facpl.sf.net/eHealth/demo.html` shows a proof-of-concept demo on how a FACPL-based access control system can be exploited for providing e-Health services.

In the rest of this section, we detail the FACPL Java library and IDE, while Section 9.4 reports performance and functionality comparisons with other similar tools.

8.1 The FACPL Library

The Java library we provide aims at representing and evaluating FACPL policies, hence at fully implementing the evaluation process formalised in Section 5. To this aim, driven by the formal semantics, we have defined a conformance test-suite that systematically verifies each library unit (e.g., expressions and combining algorithms) with respect to its formal specification.

For each element of the language the library contains an abstract class that provides its evaluation method. In practice, a FACPL policy is translated into a Java class that instantiates the corresponding abstract one and adds, by means of specific methods (e.g., `addObligation`), its forming elements. Similarly, a request corresponds to a Java class containing the request attributes and a reference to a context handler that can be used to dynamically retrieve additional attributes at evaluation-time.

Evaluating requests amounts to invoke the evaluation method of a policy, which coordinates the evaluation of its enclosed elements in compliance with its formal specification. In addition to the authorisation process, the library supports the enforcement process by defining the three enforcement algorithms and a minimal set of pre-defined PEP actions, i.e. `log`, `mailTo` and `compress`. Additional actions can be dynamically introduced by providing their implementation classes to the PEP initialisation method.

By way of example, we report in the following listings (an excerpt of) the Java code of Policy (P1) introduced in Section 4.3. Besides the specific methods used for adding policy elements, the Java code highlights the use of class references for selecting expression operators and combining algorithms. This design choice, together with the use of best-practices of object-oriented programming, allows the library to be easily extended with, e.g., new expression operators, combining algorithms and enforcement actions. Notice that rules are private inner classes, because they cannot be referred outside the enclosing policy sets.

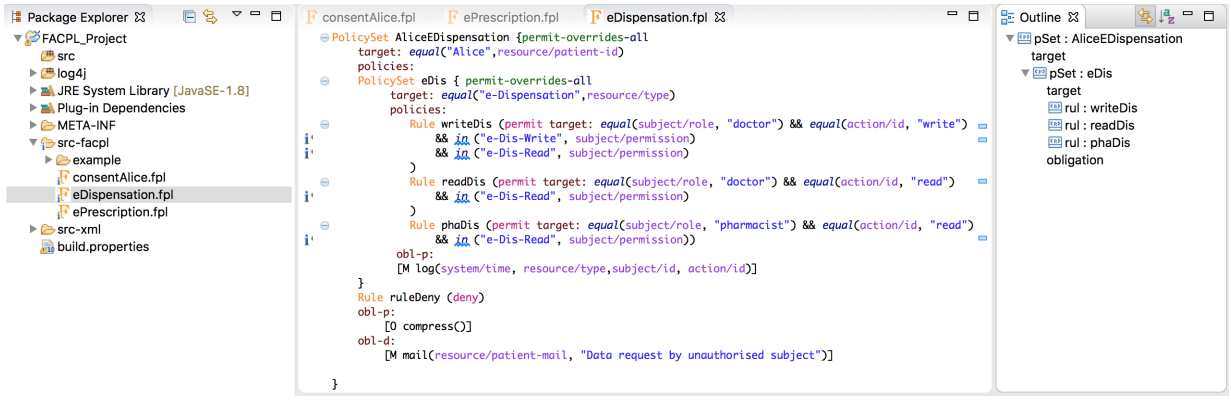

Fig. 4: The FACPL IDE

```
public class PolicySet_e-Prescription extends PolicySet{
  public PolicySet_e-Prescription(){
    addCombiningAlg(PermitOverrides.class);
    addTarget(new ExpressionFunction(Equal.class, "e-Prescription",
      new AttributeName("resource","type")));
    addRule(new rule1());
    addRule(new rule2());
    addRule(new rule3());
    addObligation(new Obligation("log",Effect.PERMIT,ObligationType.M,
       new AttributeName("system","time"),new AttributeName("resource","type"),
       new AttributeName("subject","id"),new AttributeName("action","id")));}
  private class rule1 extends Rule{
    rule1 (){
     addEffect(Effect.PERMIT);
     addTarget(...new ExpressionFunction(In.class,
    new AttributeName("subject","permission"),"e-Pre-Write"),...);}}
  private class rule2 extends Rule{ rule2 (){...} }
  private class rule3 extends Rule{ rule3 (){...} }}
```

Besides the four-valued decisions considered so far, the FACPL library also supports the extended indeterminate values used by XACML, i.e. indetP, indetD and indetDP. They can be used to specify the potential decision (permit, deny and both, respectively) that should have been returned by the evaluation of a policy if an error would not have occurred. Extended indeterminate values allow the PDP to obtain additional information about policy evaluation, which can be exploited, e.g., during policy debugging for improving the treatment of errors. However, their usage may require additional workload. In fact, it establishes that if the target of a policy set evaluates to error, rather than stopping and returning indet, the evaluation process continues the computation by processing the enclosed policies and using the decision resulting from the application of the combining algorithm to calculate an extended indeterminate value. Thus, e.g., if the combining algorithm returns the decision permit, the evaluation of the policy returns indetP. For all these reasons, we have chosen to support the extended indeterminate values by means of a boolean parameter (of the method `doAuthorisation`) whose setting can enable or disable their use at each PDP invocation.

8.2 The FACPL IDE

The FACPL IDE (see the screenshot in Figure 4) is developed as an Eclipse plug-in and aims at bringing together the available functionalities and tools. Indeed, it fully supports writing, evaluating and analysing FACPL specifications. The plug-in has been implemented by means of Xtext (`http://www.eclipse.org/Xtext/`), that is a framework to design and deploy domain-specific languages.

The plug-in accepts an enriched version of the FACPL language, which contains high level features facilitating the coding tasks. In particular, each policy has an identifier that can be used as a reference to include the policy within other policies, while specific linguistic handles enable the definition of new expression operators and combining algorithms. In order to ease the organisation of large policy specifications, the plug-

in supports modularisation of files and import commands extending file scopes.

The development environment provided by the plug-in is standard. It offers graphical features (e.g., keywords highlighting, code suggestion and navigation within and among files), static controls on FACPL code (e.g., uniqueness of identifiers and type checking), and automatic generation of Java, XACML, and SMT-LIB code. To configure all the required libraries, a dedicated wizard creates a FACPL-type project.

To facilitate the analysis of FACPL policies, the plug-in also provides a simple interface allowing policy developers to specify the authorisation and structural properties to be verified on a certain policy. Thus, the plug-in automatically generates the corresponding SMT-LIB files according to the strategies reported in Section 7.4; an execution script for the Z3 solver is also generated. Of course, the SMT-LIB files can be also evaluated by any other solver accepting SMT-LIB and supporting the theories we use.

As previously pointed out, the Java library is flexible enough to be easily extended. The plug-in facilitates this task by means of dedicated commands. For instance, to define a new expression operator, once a developer has defined the signature of the new function (which is used for type checking and inference), a template of its Java and SMT-LIB implementation is automatically generated. The actual implementation of the Java class, as well as that of the SMT-LIB function, is left to the developer.

In the FACPL IDE, all the mentioned functionalities are offered via the customised FACPL project created by the dedicated wizard. Once the project is created, the policy developer can write code starting from the basic FACPL and XACML examples already provided or from scratch. A FACPL file is a generic text file with extension *.fpl*, which has dedicated text editor, outline view and contextual menus. Functionalities supporting code development, e.g. code suggestion and auto-completion, are available via the usual Eclipse shortcuts and menus. In particular, from either the toolbar menu or the right-click editor menu, the developer can find a set of pre-defined commands to generate Java, XACML and SMT-LIB code, or to open a step-by-step wizard for the definition of authorisation and structural properties.

9 Related Work

A preliminary version of FACPL was introduced in [36] with the aim of formalising the semantics of XACML. The language presented here addresses a wider range of aspects concerning access control. Specifically, the syntax of the language is cleaned up and streamlined (e.g., rule conditions are integrated with rule targets and the policy structure is simplified); at the same time, it is extended with additional combining algorithms, the PEP specification, an explicit syntax for expressions, and obligations. This latter extension widens FACPL applicability range and expressiveness, as it provides the policy evaluation process with further, powerful means to affect the behaviour of controlled systems (see e.g. [32] for a practical example of a policy-based manager for a Cloud platform). Additional important differences concern the definition of the policy semantics: in [36] it is given in terms of partitions of the set of all possible requests, while here it is defined in a functional fashion with respect to a generic request. The new approach also features the formalisation of combining algorithms in terms of binary operators and instantiation strategies, and the automatic management of missing attributes and evaluation errors throughout the evaluation process. Most of all, the aim of this work is significantly different: we do not only propose a different language, but we provide a complete methodology that encompasses all phases of policy lifecycle, i.e. specification, analysis and enforcement. Concerning the analysis, we define a set of relevant authorisation and structural properties (whose preliminary definition is given in [34]) characterised in terms of sets of requests. We then introduce a constraint-based representation of policies and an SMT-based approach for mechanically verifying properties on top of constraints. To effectively support the functionalities, we provide a fully-integrated software toolchain.

In the rest of this section we survey more closely related work. First, we comment on differences and interoperability of FACPL with the already mentioned standard XACML (Section 9.1). Then, we discuss other relevant policy languages (Section 9.2), and approaches to the analysis of (access control) policies (Section 9.3). Finally, we compare supporting tools (Section 9.4).

Tab. 9: FACPL vs. XACML on the e-Health case study

Policy	Number of lines		Saved	Number of types		Saved
	XACML	FACPL	lines	XACML	FACPL	types
e-Prescription	239	24	89,95%	10.656	894	91,61%
e-Dispensation	239	24	89,95%	10.674	914	91,43%
Patient Consent	423	38	91,01%	19.195	1.558	91,88%

9.1 FACPL vs XACML

XACML [38] is a well-established standard for the specification of attribute-based access control policies and requests. It has an XML-based syntax and an evaluation process defined in accordance with [49] (hence similar to the FACPL one). As a matter of notation, hereafter the words emphasised in sans-serif, e.g. Rule, are XML elements, while element attributes are in italics.

From a merely lexical point of view, FACPL allows developers to define each policy element via a lightweight mnemonic syntax and leads to compact policy specifications. Instead, the XML-based syntax used by XACML ensures cross-platform interoperability, but generates verbose specifications that are hardly of immediate comprehension for developers and are not suitable for formally defining semantics and analysis techniques. Table 9 exemplifies a lexical comparison between the FACPL policies for the e-Health case study and the corresponding XACML ones (both groups of policies can be downloaded from `http://facpl.sf.net/eHealth/`).

Although FACPL and XACML policies have a similar structure, there are quite a number of (semantic) differences. Hereafter we outline the main ones.

In FACPL, request attributes are referred by structured names. In XACML, they are referred by either AttributeDesignator or AttributeSelector elements. The former one corresponds to a typed version of a structured name, while the latter one is defined in terms of XPath expressions, which are not supported by FACPL. Anyway, FACPL can represent some of them by appropriately using structured names; e.g. an AttributeSelector with category *subject* and an XPath expression like $type/id/text()$ correspond to subject/type.id.

A XACML Target is made of Match elements defining basic comparison functions on request attributes. The elements are then organised in terms of the tag structure AnyOf-AllOf-Match. This structure can be rendered in FACPL by means of, respectively, the expression operators and-or-and. However, slightly different results can be obtained from target evaluations due to the management of errors and missing attributes. Indeed, when a value is missing, the XACML semantics of Match elements returns error or false (according to the setting of the boolean parameter MustBePresent), whereas the FACPL semantics of the target elements (depending on the expression operator) can return $\perp$ possibly until the level of policies is reached, where $\perp$ is converted to false; the same occurs for evaluation errors. From our point of view, the FACPL management of missing attributes and evaluation errors is smoother than the XACML's one. Indeed FACPL can distinguish if a boolean target function has returned false due to a not satisfied comparison or due to a missing attribute; instead, XACML cannot always do it.

Additionally, the evaluation of Match functions in XACML is iteratively defined on all the retrieved attribute values. To ensure a similar behaviour in FACPL, a XACML expression such as, e.g., an equality comparison must be translated into an operator defined on sets, like e.g. in. Clearly, this limits the amount of XACML functions that can be faithfully represented in FACPL.

Except for target evaluation, the semantics of XACML and FACPL policies mainly comply with each other. However, the specification approach fostered by FACPL is more generic and poses less constraints on the policy structure. In particular, XACML prescribes a policy structure based on Policys, i.e. collections of Rules, and PolicySets, i.e. collections of Policys and/or PolicySets. Most of all, XACML forces specific constraints on targets of Policy and Policy Set: they can only contain comparison functions and each comparison can only contain one attribute name. Moreover, XACML supports fewer combining algorithms than FACPL, as well as instantiation strategies (indeed, XACML only provides the greedy one). Additionally, as previously pointed out, XACML specialises the decision indet into three sub-decisions: for the sake of presentation, we

Tab. 10: Comparison of some relevant policy languages (✓* means that user encoding is required)

Features	**XACML**	**Ponder**	**ASL**	**PTaCL**	[41]	[2]	**FACPL**
Rule-based	✓	✓					✓
Logic-based			✓	✓	✓	✓	
Mnemonic spec.		✓					✓
Comb. algorithms	✓		✓*	✓*	✓	✓*	✓
Obligations	✓	✓					✓
Missing attributes	✓			✓			✓
Error handling	✓						✓

have not considered them in the formal development of FACPL, but they are fully supported by the FACPL library (see Section 8).

FACPL and XACML share the same management of obligations, although in FACPL this process is specified in a more precise manner, i.e. through the instantiation strategies and the binary combining operators. It is also worth noticing that the XACML 'obligation fulfilment' is termed 'obligation instantiation' in FACPL, since indeed the evaluation of obligations by the PDP does not carry out any task beyond its mere instantiation.

Also, XACML provides some constructs that do not crucially affect policy expressiveness and evaluation. For instance, Variable elements permit defining pointers to expression declarations. These constructs are not directly supported by FACPL.

To sum up, except for minor differences on tangled XACML aspects mainly concerning the management of missing attributes and evaluation errors, FACPL subsumes XACML policies not containing XML raw data by offering, at the same time, a higher flexibility in the policy specification approach and a richer set of combining algorithms. Most differently from XACML, FACPL provides a formal semantics that supports formally-based analysis techniques.

9.2 Policy Languages for Access Control

Policy languages have recently been the subject of extensive research, both by industry and academia. Indeed, policies permit managing different important aspects of system behaviours, ranging from access control to adaptation and emergency handling. We compare in the following the main policy languages devoted to access control, which is our focus; Table 10 summarises the comparison.

Among the many proposed policy languages, we can identify two main specification approaches: *rule-based*, as e.g. the XACML standard and Ponder [11, 47], and *logic-based*, as e.g. ASL [24], PTaCL [8] and the logical frameworks in [2]. Many other works, as e.g. [31, 41, 39], study (part of) XACML by formally addressing peculiar features of design and evaluation of access control policies.

In the rule-based approach, policies are structured into sets of declarative rules. The seminal work [42] introduces two types of policies: authorisations and obligations. Policies of the former type have the aim of establishing if an access can be performed, while those of the latter type are basically Event-Condition-Action rules triggering the enforcement of adaptation actions. This setting is at the basis of Ponder.

Ponder is a strongly-typed policy language that, differently from FACPL, takes authorisation and obligations policies apart. Ponder does not provide explicit strategies to resolve conflictual decisions possibly arising in policy evaluation, rather it relies on abductive reasoning to statically prevent conflicts from occurring, although no implementation or experimental results are presented. On the contrary, FACPL provides combining algorithms, as we think they offer higher degrees of freedom to policy developers for managing conflicts. Similarly to Ponder, FACPL uses a mnemonic textual specification language and addresses value types, although they are not explicitly reported. Finally, the FACPL evaluation process is triggered by requests and not by events as in Ponder. Anyway, the FACPL approach is as general as the Ponder one since, by exploiting attributes, requests can represent any event of a system.

The logic-based approach mainly exploits predicate or multi-valued logics. Most of these proposals are

based on Datalog [7] (see, e.g., [24, 20, 14]), which implies that the access rules are defined as first order logic predicates. In general, these approaches offer valuable means for a low-level design of rules, but the lack of high-level features, e.g. combining algorithms or obligations, prevent them from representing policies like those of FACPL.

ASL is one of the firstly defined logic-based languages. It expresses authorisation policies based on user identity credentials and authorisation privileges, and supports hierarchisation and propagation of access rights among roles and groups of users. Additional predicates enable the definition of (a posteriori) integrity checks on authorisation decisions, e.g. conflict resolution strategies. Differently from ASL, FACPL provides high-level constructs and offers by-construction many not straightforward features like, e.g., conflict resolution strategies. A suitable use of policies hierarchisation enables propagation of access rights also in FACPL specifications.

PTaCL follows the logic-based approach as well, but it does not rely on Datalog. It defines two sets of algebraic operators based on a multi-valued logic: one modelling target expressions, the other one defining policy combinations. These operators emphasise the role of missing attributes in policy evaluation, in a way similar to FACPL, but only partially address errors. In fact, combination operators are not defined on error values: it is rather assumed that all target functions are string equalities that never produce errors. Similarly to FACPL, PTaCL permits formalising the *non-monotonicity* and *safety* properties of attribute-based policies introduced in [45]. The PTaCL extension reported in [10] introduces obligations and their instantiation, but it still lacks error handling.

A similar study, but more focussed on the distinguishing features of XACML, is reported in [39]. It introduces a formalisation of XACML in terms of multi-valued logics, by first considering 4-valued decisions and then 6-valued ones. Most of the XACML combining algorithms are formalised as operators on a partially ordered set of decisions, while the algorithms first-app and one-app are defined by case analysis. Differently from FACPL, this formalisation does not deal with missing attributes and obligations, which have instead a crucial role in XACML policy evaluation.

Another logic-based language is presented in [2]. In this case, a policy is a list of constraint assertions that are evaluated by means of an SMT solver. The framework supports reasoning about different properties, but any high-level feature, e.g. combining algorithms, has to be encoded 'by hand' into low-level assertions. In addition, missing attributes, erroneous values and obligations are not addressed.

Multi-valued logics and the relative operators have also been exploited to model the behaviour of combining algorithms. For example, the *Fine-Integration Algebra* introduced in [41] models the strategies of XACML combining algorithms by means of a set of 3-valued (i.e., permit, deny and not-app) binary operators. The behaviour of each algorithm is then defined in terms of the iterative application of the operators to the policies of the input sequence. This approach significantly differs from the FACPL one since it does not consider the indet decision. Instead, [31] explicitly introduces an error handling function that, given two decisions, determines whether their combination produces an error, i.e. an indet decision. Each (binary) operator is then defined using such error function. The formalisation of FACPL combining algorithms follows a similar approach, but it also deals with obligations and instantiation strategies, which require different iterative applications of operators.

Moreover, in [31] nonlinear constraints are used for the specification of combining algorithms which return a decision *dec* if the majority of the input policies return *dec*. Such algorithms are not usually dealt with in the literature and cannot be expressed in terms of iterative application of some binary operators.

9.3 Analysis of Access Control Policies

The increasing spread of policy-based specifications has prompted the development of many verification techniques like, e.g., property checking and behavioural characterisations. Such techniques have been implemented by means of different formalisms, ranging from SMT formulae to multi-terminal binary decision diagrams (MTBDD), including different kinds of logics. Hereafter we review the more relevant ones.

The works concerning policy analysis that are closer to our approach are of course those exploiting SMT formulae. In [46], a strategy for representing XACML policies in terms of SMT formulae is introduced. The representation, which is based on an informal semantics of XACML, supports integers, booleans and

reals, while the representation of sets of values and strings is only sketched. The combining algorithms are modelled as conjunctions and disjunctions of formulae representing the policies to be combined, i.e. in a form similar to the approach shown in Appendix A. As a design choice, formulae corresponding to the not-app decision are not generated, because they can be inferred as the complementary of the other ones. Thus, in case of algorithms like d-unless-p, additional workload is required. Moreover, the representation assumes that each attribute name is assigned only to those values that match the implicit type of the attribute, hence the analysis cannot deal with missing attributes or erroneous values. Finally, it does not take into account obligations, which have instead an important role in the evaluation. The SMT-based framework of [2], mentioned in Section 9.2, suffers from similar drawbacks.

The only analysis approach that takes missing attributes into account is presented in [9]. The analysis is based on a notion of request extension, as we have done in Section 7. Differently from our approach, this analysis aims at quantifying the impact of possibly missing attributes on policy evaluations.

The change-impact analysis of XACML policies presented in [18] aims at studying the consequences of policy modifications. In particular, to verify structural properties among policies by means of automatic tools, this approach relies on an MTBDD-based representation of policies. However, it cannot deal with many of the XACML combining algorithms and, as outlined in [2], an SMT-based approach like ours scales significantly better than the MTBDD one.

Datalog-based languages, like e.g. ASL, only provide limited analysis functionalities, that are anyway significantly less performant than SMT-based approaches. In general, these languages are useful to reason on access control issues at an high abstraction level, but they neglect many of the advanced features of modern access control systems.

Description Logic (DL) is used in [26] as a target formalism for representing a part of XACML. The approach does not take into account many combining algorithms and the decisions not-app and indet. Thus, it only permits reasoning on a set of properties significantly reduced with respect to that supported by our SMT-based approach. Furthermore, DL reasoners support the verification of structural properties of policies but suffer from the same scalability issues as the MTBDD-based reasoners.

Answer Set Programming (ASP) is used in [1, 40] for encoding XACML and enabling verification of structural properties that are similar to the `complete` one defined in Section 7.1.2. This approach however suffers from some drawbacks due to the nature of ASP. In fact, differently from SMT, ASP does not support quantifiers and multiple theories like datatype and arithmetic. Some seminal extensions of ASP to "Modulo Theories" have been proposed, but, to the best of our knowledge, no effective solver like Z3 is available. Similarly, the work in [22] exploits the SAT-based tool Alloy [23] to detect inconsistencies in XACML policies. However, as outlined in [2] and [18], Alloy is not able to manage even quite small policies and, more importantly, it cannot reason on arithmetic or any additional theory.

Finally, it is worth noticing that various analysis approaches using SAT-based tools have been developed for the Ponder language, see e.g. [3]. These approaches, however, cannot actually be compared with ours due to the numerous differences among Ponder and FACPL. Furthermore, many other works deal with the analysis of access control policies by using, e.g., process algebra and model checking techniques. However, they consider only a limited part of access control policy aspects and suffer from scalability issues with respect to SMT-based tools.

In summary, all the approaches to the analysis of access control policies mentioned above are deficient in several respects when compared with ours. Those based on SMT formulae do not address relevant aspects like, e.g. missing attributes, while the other ones do not enjoy the benefits of using SMT, i.e. support of multiple theories and scalable performance.

9.4 Supporting Tools: Performance and Functionalities

The effectiveness of supporting tools is a crucial point for the usability of a policy language. Therefore, hereafter we examine the performance of both the FACPL Java library and the SMT-based automatic analysis, and the functionalities offered by the FACPL IDE.

Concerning the library, we conducted two different tests[12]: (i) a performance comparison with a state-of-the-art XACML tool on the CONTINUE [28] case study (partially analysed in [18]); (ii) a performance stress test on a large set of randomly generated policies, thus to analyse the scalability of the library. We present below our test results, focussing on the most relevant ones; the suites of policies and requests, as well as all test results, are available at `http://facpl.sf.net/test/`.

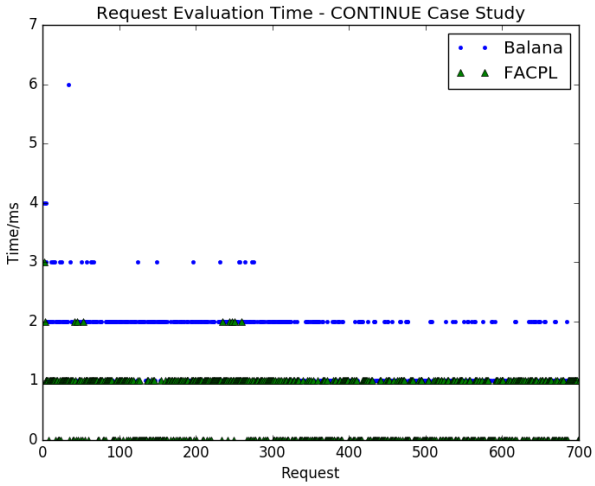

Fig. 5: FACPL vs. Balana performance evaluation

The XACML standard is by now the point-of-reference for industrial access control. In the authors' knowledge, the most up-to-date, freely available XACML tool is Balana [48]. Differently from our framework that represents FACPL policies as Java classes, Balana manages XACML policies directly in XML by exploiting a DOM representation of the XML files and evaluating XACML requests through a visit of the DOM representing the policy. We have compared the evaluation of more than 1.500 requests and obtained the results reported in Figure 5; for the sake of readability only 700 requests are reported. The mean request evaluation time is 0.49ms for FACPL and 1.27ms for Balana: evaluating a Java class ensures higher performance than navigating the DOM. Additionally, Balana requires an initial set-up time of 770ms to create the DOM.

The CONTINUE case study is by now used as a standard benchmark in the field of access control tools. However it is relatively small: it is made of 24 policies controlling 14 attributes. All policies are combined through the first-app algorithm thus, as soon as a policy applies, the evaluation stops. Therefore, for evaluating performance and scalability of the FACPL Java library, we have also considered a set of large randomly-generated policies. We generated the policies according to the following criteria: (i) a variable number of occurring *attribute names* (i.e., 10, 100, 1.000 or 10.000); (ii) a variable policy *depth* (i.e., from

[12] Both tests were conducted on a MacBook Pro, 3.1 GHz Intel i7, 16 Gb RAM, running OS X Sierra.

$p(d,w,a)$				
$\#w^1$	p_1^1	$\cdots$	p_w^1	
$\#w^2$	$p_1^2 \ldots p_w^2$	$\cdots$	$p_{w+1}^2 \cdots p_{w^2}^2$	
$\vdots$	$\vdots$	$\vdots$	$\vdots$	
$\#w^d$	$p_1^d \ldots p_w^d$	$\cdots$	$p_{w^d-w+1}^d \cdots p_{w^d}^d$	

(a)

$d \backslash w$	1	2	3	4	5
1	1	2	3	4	5
2	2	6	12	20	30
3	3	14	39	84	155
4	4	30	120	340	780
5	5	62	363	1364	3905

(b)

Tab. 11: (a) Structure of each of policy $p(d,w,a)$, where a is the number of occurring attributes names; (b) Total number of sub-policies for each combination of d and w

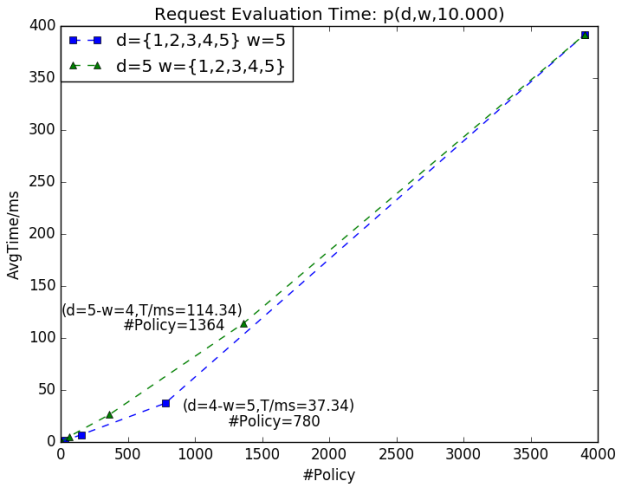

Fig. 6: FACPL performance stress test (when $a = 10.000$)

1 to 5); (iii) a variable policy *width* (i.e., from 1 to 5); (iv) only the all instantiation strategy is used (so to always require the evaluation of all the occurring policies). The combinations of the previous criteria give rise to a test-bed of 100 policies, formed by a distinct number of differently structured sub-policies featuring a different number of attribute names. More specifically, given a number of attribute names a, depth d and width w, the policy $p(d, w, a)$ is generated according to the template in Table 11a. Namely, d corresponds to the nesting levels in the policy hierarchy, while w corresponds to the number of policies that each policy (set) in the hierarchy contains. The total number of sub-policies contained by every policy $p(d, w, a)$ is summarised in Table 11b. For each of the 25 generated policies, there are 4 different versions, one for each value that a can take.

The generated test-bed has been used to perform the stress test on the FACPL library. The results, when a is set to 10.000, are summarised in Figure 6. The graphs show how the performance changes as a function of the policy structure, i.e. depending on d and w. Better performances are obtained by structuring policies in terms of larger width values (marked by the blue square), rather than larger depth values (marked by the green triangle). Namely, the average evaluation time increases more significantly by increasing policy depth than width.

Concerning the automatic analysis, as previously pointed out, the tool closer to ours is that of [2], which relies on the SMT solver Yices [15]. Differently from Z3, Yices does not support datatype theory, which is instead crucial to deal with a wide range of policy aspects, as e.g. missing and erroneous attributes. To analyse the completeness of the CONTINUE policies, the Yices-based tool requires around 570ms, while our Z3-based tool requires around 120ms[13]. Notably, other tools not based on SMT, like e.g. Margrave, have significantly worse performance when policies scale. In fact, as reported in [2], the increment of the number of possible values for the attributes occurring in the CONTINUE policies prevents Margrave to accomplish the analysis. On the contrary, SMT solvers can also deal with infinite sets of attribute values, as e.g. integers. To further evaluate the analysis performance, we also report in Figure 7 the time required to verify the satisfiability of the not-app constraint (i.e., the verification of the `complete` property; marked by the blue square), and the permit constraint (marked by the green triangle) of the policies $p(5, 5, a)$, i.e. 3905 policies, with a varying number a of attributes. Namely, the `complete` property is always verified in less than one second, despite the increasing number of attributes. Instead, the verification time for the satisfiability of the permit constraint increases by rising the number of attribute names, but the increments are significantly lower than the attribute name variations, i.e. from $199s$ with 1.000 attribute names, to $394s$

[13] The Yices value is taken directly from [2], since the provided CONTINUE implementation only runs on Windows machines. Therefore, we ran this Z3 analysis on an older comparable hardware configuration (with the current configuration it takes only 60ms).

with 10.000 ones. The difference between the two cases is due to the policy semantics: as soon as a policy target is not-app, the whole policy is not-app, while a policy evaluates to permit according to the combination strategies of the chosen combining algorithms. It is finally worth noticing that the considered policies have a limited number of attribute names representing set values. In fact, a higher number of set attributes would require many satisfiability checks of existential quantifiers along with the 3905 policies (see the definition of the `inInt` function in Section 7.3). In such a case, the analysis remains feasible, i.e. hours instead of minutes, according to the cases. For example, it lasts two hours for the case of 100 attribute names and 20 set attributes, while it cannot complete for the case of 1.000 attribute names and 200 set ones.

We conclude by commenting on the most strictly related IDEs. To the best of our knowledge, the only freely available IDEs are the ALFA Eclipse plugin by Axiomatics (`http://www.axiomatics.com/alfa-plugin-for-eclipse.html`) and the graphical editor of the Balana-based framework (`http://xacmlinfo.org/category/xacml-editor/`). However, differently from our IDE, they only provide a high-level language for writing XACML policies. Additionally, ALFA does not provide any request evaluation engine, since the Axiomatics one is a proprietary software.

10 Concluding Remarks and Future Work

We have described a full-fledged framework for the specification, analysis and enforcement of access control policies. Our framework relies on FACPL and is built on top of solid formal foundations. The FACPL semantics provides a formalisation of complex access control features —including obligations and missing attributes, which are instead overlooked by many other proposals— and lays the basis for developing analysis techniques and tools. We have shown that FACPL policies can be represented in terms of specific SMT formulae, whose automatic evaluation permits verifying various authorisation and structural properties. We have demonstrated feasibility and effectiveness of our approach by means of a case study from the e-Health application domain, which is currently one of the most critical application domains of access control systems. We have also shown that the use of SMT solvers provides us with stable and efficient tools, ensuring better performance than many other approaches from the literature.

In a general perspective, our approach brings together the benefits deriving from using a high-level, mnemonic rule-based language with the rigorous means provided by denotational semantics and constraints. Most of all, the supporting tools we implemented allow access control system developers to use any of the formally-defined functionalities provided by our framework, without the need that they be familiar with formal methods.

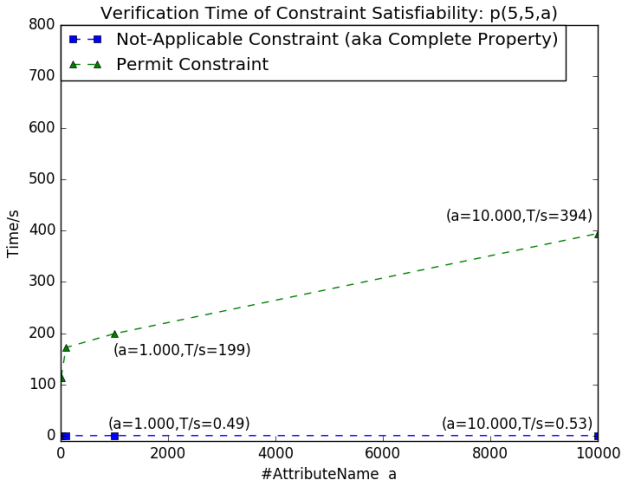

Fig. 7: FACPL analysis performance (when $d = w = 5$)

We conclude by first enlighting some distinguishing traits of FACPL (Section 10.1), then by pointing out some future research directions (Section 10.2).

10.1 Discussion

We want here to recap and reflect on a few characteristics of FACPL (and its framework) and the design choices that underlie them.

Expressiveness. The access control systems expressible by FACPL are those expressible by XACML (but not dealing with XML raw data, see Section 9.1), with in addition the possibility of using consensus-based combining algorithms and the all instantiation strategy for obligations. FACPL access control systems are systematically more compact (see Table 9) and feature a smoother management of errors and missing attributes. This latter characteristic, together with the fact that we decided not to introduce any static check on the type of requests, permits to accurately deal with every access control request. Alternatively, we could have defined a type inference system in order to statically check the policies and infer the expected type of any attribute name occurring within. Then, we could have reserved evaluation for only those requests whose attribute names comply with their expected type, while we could have directly returned the indet decision for all the other requests. By pursuing such an approach, however, we would have lost expressiveness, since policies could not be able anymore to automatically manage errors due to unexpected attribute values and possibly mask them by using operators that combine, according to different strategies, indet decisions with the others.

Besides the definition of access controls, FACPL permits defining obligations, which are a key ingredient to enhance the expensiveness of access control systems. As exemplified in the definition of the e-Health case study (see Section 4.3), the instantiation of obligations permits defining context-dependent actions to be enforced at run-time in the controlled system. Indeed, the side-effects of policy evaluation are not only the enforcement of access decisions, but also the enforcement of dynamically instantiated actions. FACPL obligations permit enforcing, e.g., resource usage, adaptation and emergency handling strategies. For example, in the application of the Ponder language of [43] and in the preliminary version of FACPL of [33, 32], obligations are used to enforce self-adaptation strategies in autonomic computing systems. Instead, in the context of emergency handling, an obligation-based approach is proposed in [6, 35] to enforce the principle known as 'break the glass', which means that authorisation controls can be bypassed in case of emergency.

We intentionally abstract from the actual syntax of obligations. They are simply intended to be actions executed at run-time. From time to time, they can be chosen to more adequately express the access control system in hand. We also abstract from the actual semantics obligations. Indeed, the discharging of obligations done by the PEP simply refers to the fact that the system has taken charge of their execution, which is intended to finish by the conclusion of the PEP enforcement process. However, the possibility of enforcing some obligations after releasing the decision and granting the access is a topic worth to be studied (it is indeed one of the future research directions we want to pursue).

Validation. Our framework has essentially three constituent elements: (i) the linguistic constructs together with their denotational semantics; (ii) the constraint formalism and the semantic-preserving translation; (iii) the Java-based supporting tools. For each of them we have presented different validation results, both theoretical and empirical.

The linguistic constructs are validated with respect to their expressiveness. This is done, on the one hand, by modelling a real-world case study (Section 4.3) from the e-Health application domain, on the other hand, by comparing (Section 9.1) FACPL with XACML, i.e. the state-of-the-art OASIS standard for attribute-based access control systems. FACPL formal semantics is validated according to the so-called *reasonability properties* of [45] that precisely characterise the expressiveness of a policy language. Besides these properties, we show that the semantics is well-defined (Theorem 5.1) and precisely characterise the attributes that are relevant for policy evaluation (Lemma 5.2); this important result, as pointed out in Section 7.1.3, underlies the automatic property verification. All the results are presented in Section 5.6, while their proofs are relegated to Appendix B.1.

Similarly, the constraint formalism and the semantic-preserving translation of FACPL policies into SMT

formulae are validated by the theoretical results presented in Section 6.3 and proved in Appendix B.2. All together these results ensure that the approach to the analysis of FACPL policies presented in Section 7 is sound.

The software tools are validated by empirically examining their performance and functionalities. The obtained results are reported in Section 9.4.

Exploitation. The FACPL framework is a production-level software that is also used in industry. Indeed, since its preliminary version, FACPL has been used by Tiani Spirit (`http://www.tiani-spirit.com/`) instead of XACML to carry out design and automatic analysis of access control policies. In particular, the FACPL access control engine has been used as XACML reference implementation in several projects. Furthermore, FACPL was used for team works in a PhD school on engineering Collective Autonomic Systems (`http://www.ascens-ist.eu/springschool`) and has been used in many bachelor and master thesis (further details can be found at the FACPL web-site). These practical exploitations have highlighted that its compact, mnemonic syntax requires very short learning time, even to undergraduate students. The users have also appreciated the flexibility of the IDE, which can be smoothly integrated within other development environments.

Extendability. The proposed framework offers a variegated set of constructs, ranging from expression operators to combining algorithms, for defining access controls. Anyway, to better suit any need, as reported in Section 8, both the Java library and the IDE can be easily extended with the introduction of, e.g., new expression operators. This approach supports writing customised FACPL specifications. These specifications can then be translated, in accordance with the user's definition of the added constructs, to Java and SMT-LIB code that can be still evaluated and analysed, respectively. The formal assurance of semantic preservation (Theorem 6.2) can be easily tailored for encompassing the user's extensions. For instance, in case of addition of new expression operators, it only requires devising a constraint operator (or a combination thereof) that faithfully represents the semantics of the new operator.

10.2 Future Work

In the next future, we plan to address the issue of controlling the accesses while they are in progress. In this sort of 'continuative' access control, the challenge is to ensure guarantees on how granted accesses are used. This model is usually referred to as Usage Control [30] in the literature and has been recently studied by various researchers. To deal with usage control, temporal aspects are of paramount importance, both to refer to ongoing accesses and to enforce obligations after releasing access decisions. To this aim, we will provide a FACPL-based solution for usage control that, by relying on the already available context-dependent authorisation process, can control ongoing accesses and instantiate temporal obligations. To actually enforce these obligations and, consequently, reason on them, we will refine the PEP semantics by appropriately instantiating the predicate $\Downarrow\mathsf{ok}$ introduced in Section 5.4.

We also plan to provide a formally-based analysis technique that system developers can exploit to verify, e.g., history-dependent properties like dynamic separation of duty. To this aim, besides formalising new history-dependent authorisation properties, we want to define and verify properties on conflicts and dependencies among obligations.

References

[1] G. J. Ahn, H. Hu, J. Lee, and Y. Meng. Representing and reasoning about web access control policies. In *COMPSAC*, pages 137–146. IEEE Computer Society, 2010.

[2] K. Arkoudas, R. Chadha, and C.-Y. J. Chiang. Sophisticated access control via SMT and logical frameworks. *ACM Trans. Inf. Syst. Secur.*, 16(4):17, 2014.

[3] A. K. Bandara, E. Lupu, and A. Russo. Using event calculus to formalise policy specification and analysis. In *POLICY*, page 26. IEEE, 2003.

[4] C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanovic, T. King, A. Reynolds, and C. Tinelli. CVC4. In *Proc. of CAV*, volume 6806 of *LNCS*, pages 171–177. Springer, 2011.

[5] A. Bertolino, S. Daoudagh, F. Lonetti, and E. Marchetti. The X-CREATE Framework - A Comparison of XACML Policy Testing Strategies. In *WEBIST*, pages 155–160. SciTePress, 2012.

[6] A. D. Brucker and H. Petritsch. Extending access control models with break-glass. In *SACMAT*, pages 197–206. ACM, 2009.

[7] S. Ceri, G. Gottlob, and T. Tanca. What you always wanted to know about datalog (and never dared to ask). *IEEE Trans. Knowl. Data Eng.*, 1(1):146–166, 1989.

[8] J. Crampton and C. Morisset. Ptacl: A language for attribute-based access control in open systems. In P. Degano and J. D. Guttman, editors, *POST*, volume 7215 of *LNCS*, pages 390–409. Springer, 2012.

[9] J. Crampton, C. Morisset, and N. Zannone. On missing attributes in access control: Non-deterministic and probabilistic attribute retrieval. In *SACMAT*, pages 99–109. ACM, 2015.

[10] J. Crampton and C. Williams. Obligations in ptacl. In Sara Foresti, editor, *STM*, volume 9331 of *LNCS*, pages 220–235. Springer, 2015.

[11] N. Damianou, N. Dulay, E. Lupu, and M. Sloman. The Ponder Policy Specification Language. In *POLICY*, LNCS 1995, pages 18–38. Springer, 2001.

[12] L. M. de Moura and N. Bjørner. Z3: an efficient SMT solver. In *TACAS*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.

[13] L. M. de Moura and N. Bjørner. Satisfiability modulo theories: introduction and applications. *Commun. ACM*, 54(9):69–77, 2011.

[14] J. DeTreville. Binder, a logic-based security language. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, SP '02, pages 105–113, Washington, DC, USA, 2002. IEEE Computer Society.

[15] B. Dutertre. Yices 2.2. In *Proc. of CAV*, volume 8559 of *LNCS*, pages 737–744. Springer, 2014.

[16] European Parliament and Council. Directive 95/46/EC, 1995. Official Journal L 281 , 23/11/1995 P. 0031 - 0050. `http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:en:HTML`.

[17] D. F. Ferraiolo and D. R. Kuhn. Role-based access control. In *NIST-NCSC National Computer Security Conference*, pages 554–563, 1992.

[18] K. Fisler, S. Krishnamurthi, L.A. Meyerovich, and M.C. Tschantz. Verification and change-impact analysis of access-control policies. In *ICSE*, pages 196–205. ACM, 2005.

[19] W. Han and C. Lei. A survey on policy languages in network and security management. *Computer Networks*, 56(1):477–489, 2012.

[20] M. Hashimoto, M. Kim, H. Tsuji, and H. Tanaka. Policy description language for dynamic access control models. In *DASC*, pages 37–42. IEEE, 2009.

[21] V. C. Hu, D. R. Kuhn, and D. F. Ferraiolo. Attribute-based access control. *IEEE Computer*, 48(2):85–88, 2015.

[22] G. Hughes and T. Bultan. Automated verification of access control policies using a sat solver. *STTT*, 10(6):503–520, 2008.

[23] D. Jackson. Alloy: a lightweight object modelling notation. *ACM Trans. Softw. Eng. Methodol.*, 11(2):256–290, 2002.

[24] S. Jajodia, P. Samarati, and V. S. Subrahmanian. A logical language for expressing authorizations. In *Symposium On Security And Privacy*, pages 31–42. IEEE, 1997.

[25] X. Jin, R. Krishnan, and R. S. Sandhu. A unified attribute-based access control model covering dac, MAC and RBAC. In *DBSec*, pages 41–55. Springer, 2012.

[26] V. Kolovski, J. A. Hendler, and B. Parsia. Analyzing web access control policies. In *WWW*, pages 677–686. ACM, 2007.

[27] M. Kovac. E-health demystified: An e-government showcase. *IEEE Computer*, 47(10):34–42, 2014.

[28] S. Krishnamurthi. The CONTINUE server (or, how I administered PADL 2002 and 2003). In *PADL*, volume 2562 of *LNCS*, pages 2–16. Springer, 2003.

[29] B. W. Lampson. Protection. *Operating Systems Review*, 8(1):18–24, 1974.

[30] A. Lazouski, F. Martinelli, and P. Mori. Usage control in computer security: A survey. *Computer Science Review*, 4(2):81–99, 2010.

[31] N. Li, Q. Wang, W. H. Qardaji, E. Bertino, P. Rao, J. Lobo, and D. Lin. Access control policy combining: theory meets practice. In *SACMAT*, pages 135–144. ACM, 2009.

[32] A. Margheri, M. Masi, R. Pugliese, and F. Tiezzi. Developing and Enforcing Policies for Access Control, Resource Usage, and Adaptation. A Practical Approach. In *WSFM*, volume 8379 of *LNCS*, pages 85–105. Springer, 2013.

[33] A. Margheri, R. Pugliese, and F. Tiezzi. Linguistic abstractions for programming and policing autonomic computing systems. In *UIC/ATC*, pages 404–409. IEEE, 2013.

[34] A. Margheri, R. Pugliese, and F. Tiezzi. On Properties of Policy-Based Specifications. In *WWV*, volume 188 of *EPTCS*, pages 33–50, 2015.

[35] S. Marinovic, N. Dulay, and M. Sloman. Rumpole: An introspective break-glass access control language. *ACM TISSEC.*, 17(1):2:1–2:32, 2014.

[36] M. Masi, R. Pugliese, and F. Tiezzi. Formalisation and Implementation of the XACML Access Control Mechanism. In *ESSoS*, LNCS 7159, pages 60–74. Springer, 2012.

[37] NIST. A survey of access control models, 2009. `http://csrc.nist.gov/news_events/privilege-management-workshop/PvM-Model-Survey-Aug26-2009.pdf`.

[38] OASIS XACML TC. eXtensible Access Control Markup Language (XACML) version 3.0 , January 2013. `https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml`.

[39] C. D. P. K. Ramli, Nielson H. Riis, and F. Nielson. The logic of XACML. *Sci. Comput. Program.*, 83:80–105, 2014.

[40] C. D. P. K. Ramli, H. Riis Nielson, and F. Nielson. Xacml 3.0 in answer set programming. In *LOPSTR*, volume 7844 of *LNCS*, pages 89–105. Springer, 2012.

[41] P. Rao, D. Lin, E. Bertino, N. Li, and J. Lobo. An algebra for fine-grained integration of XACML policies. In *SACMAT*, pages 63–72. ACM, 2009.

[42] M. Sloman. Policy driven management for distributed systems. *J. Network Syst. Manage.*, 2(4):333–360, 1994.

[43] Morris Sloman and Emil C. Lupu. Engineering policy-based ubiquitous systems. *Comput. J.*, 53(7):1113–1127, 2010.

[44] The Article 29 Data Protection WP, 2013. http://ec.europa.eu/justice/data-protection/article-29/.

[45] M. C. Tschantz and S. Krishnamurthi. Towards reasonability properties for access-control policy languages. In *SACMAT*, pages 160–169. ACM, 2006.

[46] F. Turkmen, J. den Hartog, S. Ranise, and N. Zannone. Analysis of XACML policies with SMT. In *POST*, volume 9036 of *LNCS*, pages 115–134. Springer, 2015.

[47] K. P. Twidle, N. Dulay, E. Lupu, and M. Sloman. Ponder2: A policy system for autonomous pervasive environments. In *ICAS*, pages 330–335. IEEE, 2009.

[48] WSO2. Balana: Open source XACML implementation, 2015. https://github.com/wso2/balana.

[49] R. Yavatkar, D. Pendarakis, and R. Guerin. A Framework for Policy-based Admission Control. RFC 3060 (Proposed Standard), 2000.

A Definitions for Combining Algorithms

In this section we report all the definitions regarding the combining algorithms. Table 12 shows all the combination matrices defining the binary operators $\otimes$alg for each algorithm alg. Hereafter we report the constraint resulting from the combination of two constraint tuples, say A and B, defined according to the various combining algorithms.

p-over$(A, B) =$

$\langle$ permit : $A\downarrow_p \vee B\downarrow_p$
deny : $(A\downarrow_d \wedge B\downarrow_d) \vee\ (A\downarrow_d \wedge B\downarrow_n) \vee\ (A\downarrow_n \wedge B\downarrow_d)$
not-app : $A\downarrow_n \wedge B\downarrow_n$
indet : $(A\downarrow_i \wedge \neg B\downarrow_p) \vee\ (\neg A\downarrow_p \wedge B\downarrow_i)\rangle$

d-over$(A, B) =$

$\langle$ permit : $(A\downarrow_p \wedge B\downarrow_p) \vee\ (A\downarrow_p \wedge B\downarrow_n) \vee (A\downarrow_n \wedge B\downarrow_p)$
deny : $A\downarrow_d \vee B\downarrow_d$
not-app : $A\downarrow_n \wedge B\downarrow_n$
indet : $(A\downarrow_i \wedge \neg B\downarrow_d) \vee\ (\neg A\downarrow_d \wedge B\downarrow_i)\rangle$

d-unless-p$(A, B) =$

$\langle$ permit : $A\downarrow_p \vee B\downarrow_p$
deny : $\neg A\downarrow_p \wedge \neg B\downarrow_p \wedge (A\downarrow_d \vee A\downarrow_n \vee A\downarrow_i)$
$\wedge(B\downarrow_d \vee B\downarrow_n \vee B\downarrow_i)$
not-app : false
indet : false$\rangle$

p-unless-d$(A, B) =$

$\langle$ permit : $\neg A\downarrow_d \wedge \neg B\downarrow_d \wedge (A\downarrow_p \vee A\downarrow_n \vee A\downarrow_i)$
$\wedge(B\downarrow_p \vee B\downarrow_n \vee B\downarrow_i)$
deny : $A\downarrow_d \vee B\downarrow_d$
not-app : false
indet : false$\rangle$

first-app$(A, B) =$

$\langle$ permit : $A\downarrow_p \vee (B\downarrow_p \wedge A\downarrow_n)$
deny : $A\downarrow_d \vee (B\downarrow_d \wedge A\downarrow_n)$
not-app : $A\downarrow_n \wedge B\downarrow_n$
indet : $A\downarrow_i \vee (A\downarrow_n \wedge B\downarrow_i)\rangle$

one-app$(A, B) =$

$\langle$ permit : $(A\downarrow_p \wedge B\downarrow_n) \vee\ (A\downarrow_n \wedge B\downarrow_p)$
deny : $(A\downarrow_d \wedge B\downarrow_n) \vee\ (A\downarrow_n \wedge B\downarrow_d)$
not-app : $A\downarrow_n \wedge B\downarrow_n$
indet : $A\downarrow_i \vee B\downarrow_i \vee ((A\downarrow_p \vee A\downarrow_d) \wedge\ (B\downarrow_p \vee B\downarrow_d))\rangle$

weak-con(A, B)

$\langle$ permit : $(A\downarrow_p \wedge B\downarrow_p) \vee\ (A\downarrow_p \wedge \neg B\downarrow_d) \vee\ (\neg A\downarrow_d \wedge B\downarrow_p)$
deny : $(A\downarrow_d \wedge B\downarrow_d) \vee\ (A\downarrow_d \wedge \neg B\downarrow_p) \vee\ (\neg A\downarrow_p \wedge B\downarrow_d)$
not-app : $A\downarrow_n \wedge B\downarrow_n$
indet : $(A\downarrow_p \wedge B\downarrow_d) \vee\ (A\downarrow_d \wedge B\downarrow_p) \vee\ A\downarrow_i \vee B\downarrow_i$

strong-con$(A, B) =$

$\langle$ permit : $A\downarrow_p \wedge B\downarrow_p$
deny : $A\downarrow_d \wedge B\downarrow_d$
not-app : $A\downarrow_n \wedge B\downarrow_n$
indet : $A\downarrow_i \vee B\downarrow_i \vee (A\downarrow_n \wedge \neg B\downarrow_n) \vee\ (\neg A\downarrow_n \wedge B\downarrow_n)$
$\vee\ (A\downarrow_p \wedge B\downarrow_d) \vee\ (A\downarrow_d \wedge B\downarrow_p)\rangle$

⊗p-over	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
⟨permit FO_1⟩	⟨permit $FO_1 \bullet FO_2$⟩	⟨permit FO_1⟩	⟨permit FO_1⟩	⟨permit FO_1⟩
⟨deny FO_1⟩	⟨permit FO_2⟩	⟨deny $FO_1 \bullet FO_2$⟩	⟨deny FO_1⟩	indet
not-app	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
indet	⟨permit FO_2⟩	indet	indet	indet

⊗d-over	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
⟨permit FO_1⟩	⟨permit $FO_1 \bullet FO_2$⟩	⟨deny FO_2⟩	⟨permit FO_1⟩	indet
⟨deny FO_1⟩	⟨deny FO_1⟩	⟨deny $FO_1 \bullet FO_2$⟩	⟨deny FO_1⟩	⟨deny FO_1⟩
not-app	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
indet	indet	⟨deny FO_2⟩	indet	indet

⊗d-unless-p	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
⟨permit FO_1⟩	⟨permit $FO_1 \bullet FO_2$⟩	⟨permit FO_1⟩	⟨permit FO_1⟩	⟨permit FO_1⟩
⟨deny FO_1⟩	⟨permit FO_2⟩	⟨deny $FO_1 \bullet FO_2$⟩	⟨deny FO_1⟩	⟨deny FO_1⟩
not-app	⟨permit FO_2⟩	⟨deny FO_2⟩	⟨deny ϵ⟩	⟨deny ϵ⟩
indet	⟨permit FO_2⟩	⟨deny FO_2⟩	⟨deny ϵ⟩	⟨deny ϵ⟩

⊗p-unless-d	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
⟨permit FO_1⟩	⟨permit $FO_1 \bullet FO_2$⟩	⟨deny FO_2⟩	⟨permit FO_1⟩	⟨permit FO_1⟩
⟨deny FO_1⟩	⟨deny FO_1⟩	⟨deny $FO_1 \bullet FO_2$⟩	⟨deny FO_1⟩	⟨deny FO_1⟩
not-app	⟨permit FO_2⟩	⟨deny FO_2⟩	⟨permit ϵ⟩	⟨permit ϵ⟩
indet	⟨permit FO_2⟩	⟨deny FO_2⟩	⟨permit ϵ⟩	⟨permit ϵ⟩

⊗first-app	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
⟨permit FO_1⟩	⟨permit FO_1⟩	⟨permit FO_1⟩	⟨permit FO_1⟩	⟨permit FO_1⟩
⟨deny FO_1⟩	⟨deny FO_1⟩	⟨deny FO_1⟩	⟨deny FO_1⟩	⟨deny FO_1⟩
not-app	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
indet	indet	indet	indet	indet

⊗one-app	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
⟨permit FO_1⟩	indet	indet	⟨permit FO_1⟩	indet
⟨deny FO_1⟩	indet	indet	⟨deny FO_1⟩	indet
not-app	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
indet	indet	indet	indet	indet

⊗weak-con	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
⟨permit FO_1⟩	⟨permit $FO_1 \bullet FO_2$⟩	indet	⟨permit FO_1⟩	indet
⟨deny FO_1⟩	indet	⟨deny $FO_1 \bullet FO_2$⟩	⟨deny FO_1⟩	indet
not-app	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
indet	indet	indet	indet	indet

⊗strong-con	⟨permit FO_2⟩	⟨deny FO_2⟩	not-app	indet
⟨permit FO_1⟩	⟨permit $FO_1 \bullet FO_2$⟩	indet	indet	indet
⟨deny FO_1⟩	indet	⟨deny $FO_1 \bullet FO_2$⟩	indet	indet
not-app	indet	indet	not-app	indet
indet	indet	indet	indet	indet

Tab. 12: Combination matrices for the binary operators ⊗alg

B Proofs of the Results

Some of the proofs proceed by induction on the *depth* of policies, i.e. the number of their nesting levels, which is defined by induction on the syntax of policies as follows

$$\begin{array}{l} depth((e \;\; \mathsf{target} : expr \;\; \mathsf{obl} : o^*)) = 0 \\ depth(\\ \{a \;\; \mathsf{target} : expr \;\; \mathsf{policies} : p^+ \; \mathsf{obl\text{-}p} : o_p^* \;\; \mathsf{obl\text{-}d} : o_d^* \}) = \\ \qquad\qquad\qquad\qquad 1 + max(\{depth(p) \mid p \in p^+\}) \end{array}$$

Policies with depth 0 are rules, the other ones are policies containing other policies. Notationally, we will use p^i to mean that policy p has depth i and $(p^+)^i$ to mean that at least a policy in the sequence p^+ has depth i and the others have depth at most i.

B.1 Proofs of Results in Section 5

Theorem 5.1 (Total and Deterministic Semantics).

1. For all $pas \in PAS$ and $req \in Request$, there exists a $dec \in Decision$, such that $\mathcal{P}as[\![pas, req]\!] = dec$.

2. For all $pas \in PAS$, $req \in Request$ and $dec, dec' \in Decision$, it holds that
   $\mathcal{P}as[\![pas, req]\!] = dec \;\; \wedge \;\; \mathcal{P}as[\![pas, req]\!] = dec'$
   $\Rightarrow \;\; dec = dec'$.

Proof. The goal of the proof is to show that $\mathcal{P}as$ is a total and deterministic function, i.e. it is defined for all possible input pairs and always returns the same decision any time it is applied to a specific pair. If we let pas be $\{\, \mathsf{pep} : ea \;\; \mathsf{pdp} : pdp \,\}$ then, from the clause (S-8), we have that

$$\begin{array}{l} \mathcal{P}as[\![\{\, \mathsf{pep} : ea \;\; \mathsf{pdp} : pdp \,\}, req]\!] = \\ \qquad\qquad \mathcal{E}A[\![ea]\!](\mathcal{P}dp[\![pdp]\!](\mathcal{R}[\![req]\!])) \end{array}$$

Thus, since function composition preserves totality and determinism, we are left to prove that $\mathcal{R}$, $\mathcal{P}dp$ and $\mathcal{E}A$ are total and deterministic functions. Due to their inductive definition (given in Section 5), the proof proceeds by inspecting their defining clauses with the aim of checking that they satisfy the two requirements below

R1: there is one, and only one, clause that applies to each syntactic domain element (this usually follows since the definition is syntax-driven and considers all the syntactic forms that the input can assume);

R2: for each defining clause,

- the conditions in the right hand side are mutually exclusive (from the systematic use of the `otherwise` condition, it directly follows that they cover all the possible cases for the syntactic domain elements of the form occurring in the left hand side),
- the values assigned in each case of the right hand side are obtained by only using total and deterministic functions/operators/predicates.

Case $\mathcal{R}$. From its defining clauses (S-1) we get that $\mathcal{R}$ is defined on all non-empty sequences of attributes, i.e. all requests. Moreover, the conditions in the right hand side of each clause are mutually exclusive and the operator $\Cup$ is total and deterministic by definition. Thus R1 and R2 hold, which means that $\mathcal{R}$ is a total and deterministic function.

Case $\mathcal{P}dp$. To prove this case, we first prove that $\mathcal{E}$, $\mathcal{O}$, $\mathcal{A}$ and $\mathcal{P}$ are total functions.

Case $\mathcal{E}$. By an easy inspection of the clauses defining $\mathcal{E}$, an excerpt of which are in Table 5, it is not hard to believe that they satisfy R1 (since the application of the clauses is driven by the syntactic form of the input expression) and R2 above, hence $\mathcal{E}$ is a total function. Moreover, since the operator $\bullet$ is total and deterministic, from the clauses (S-2) it follows that $\mathcal{E}$ remains a total and deterministic function also when extended to sequences of expressions.

Case $\mathcal{O}$. Since $\mathcal{E}$ is a total and deterministic function also on sequences of expressions, from the clauses (S-3a) and (S-3b) it follows that R1 and R2 hold, thus $\mathcal{O}$ is a total and deterministic function both on single obligations and on sequences of obligations.

Cases $\mathcal{A}$ and $\mathcal{P}$. The definitions of $\mathcal{P}$ and $\mathcal{A}$ are syntax-driven and consider all the syntactic forms that the input can assume, thus R1 is satisfied. Now, since $\mathcal{P}$ and $\mathcal{A}$ are mutually recursive, we prove by induction on the depth of their arguments that their defining clauses satisfy R2 for all input policies.

Base Case ($i = 0$). Let us start from $\mathcal{P}$. p^0 is of the form (e target : $expr$ obl : o^*). We have hence to prove that the clause (S-4a), which is the defining clause of $\mathcal{P}$ that applies to p^0, satisfies R2. This directly follows from the fact that $\mathcal{E}$ and $\mathcal{O}$ are total and deterministic functions. Now, let us consider $\mathcal{A}$ and proceed by case analysis on a.

($a = \mathsf{alg}_{\mathsf{all}}$ for any alg). Since the clause (S-4a) satisfies (R1 and) R2, for each p_j^0 in $(p^+)^0$, $\mathcal{P}[\![p_j^0]\!]r$ is uniquely defined. Thus, since each operator $\otimes$alg is total and deterministic by construction, the clause (S-6a), to be used since the form of a, satisfies R2 (when all the input policies have depth 0).

($a = \mathsf{alg}_{\mathsf{greedy}}$ for any alg). This case is similar to the previous one, but involves the clause (S-6b). It satisfies R2 (when all the input policies have depth 0) since its conditions in the right hand side are mutually exclusive by construction (indeed, each predicate $isFinal_{\mathsf{alg}}$ and each operator $\otimes$alg is total and deterministic).

Inductive Case ($i = n + 1$). Let us start from $\mathcal{P}$. p^{n+1} is of the form $\{a$ target:$expr$ policies:$(p^+)^n$ obl-p : o_p^* obl-d : $o_d^*\}$. By the induction hypothesis, for any r, a and p_j^k in $(p^+)^n$, with $k \leq n$, the clauses defining $\mathcal{P}$ and $\mathcal{A}$ satisfy (R1 and) R2, that is $\mathcal{P}[\![p_j^k]\!]r$ and $\mathcal{A}[\![a, (p^+)^n]\!]r$ are uniquely defined. Hence, the clause (S-4b), to be used since the form of p^{n+1}, satisfies R2 as well. For $\mathcal{A}$, we can reason like in the base case by exploiting the induction hypothesis. We can thus conclude that both the clauses (S-6a) and (S-6b) satisfy R2 (for any input policy).

Therefore, $\mathcal{P}$ and $\mathcal{A}$ are total and deterministic functions.

Now, that $\mathcal{P}dp$ is a total and deterministic function directly follows from its defining clause (S-5).

Case $\mathcal{E}A$. The requirement R1 is satisfied by definition. Moreover, since the predicate $\Downarrow$ ok is total and deterministic, the same holds for the function ($\!($ $)\!$). Therefore, also R2 is satisfied by each defining clause (the conditions on $res.dec$ are trivially mutually exclusive). Hence, $\mathcal{E}A$ is a total and deterministic function.

□

Lemma 5.2 (Policy relevant attributes). For all $p \in Policy$ and $r, r' \in R$ such that $r(n) = r'(n)$ for all $n \in Names(p)$ it holds that $\mathcal{P}[\![p]\!]r = \mathcal{P}[\![p]\!]r'$.

Proof. The statement is based on an analogous result concerning expressions

$$\begin{array}{l}\text{for all } expr \in Expr \text{ and } r_1, r_1' \in R \text{ such that}\\ r_1(n) = r_1'(n) \text{ for all } n \in Names(expr),\\ \text{it holds that } \mathcal{E}[\![expr]\!]r_1 = \mathcal{E}[\![expr]\!]r_1'\end{array} \quad \text{(R)}$$

which can be easily proven by structural induction on the syntax of expressions. Functions r_1 and r_1' are only exploited in the base case when evaluating a name $n \in \mathit{Names}(\mathit{expr})$ for which, by definition and hypothesis, we have $\mathcal{E}[\![n]\!]r_1 = r_1(n) = r_1'(n) = \mathcal{E}[\![n]\!]r_1'$. Since for any *expr* occurring in p, we have that $\mathit{Names}(\mathit{expr}) \subseteq \mathit{Names}(p)$, from (R), by taking $r_1 = r$ and $r_1' = r'$, it follows that

$$\text{for all } \mathit{expr} \text{ occurring in } p,\ \mathcal{E}[\![\mathit{expr}]\!]r = \mathcal{E}[\![\mathit{expr}]\!]r' \tag{R-E}$$

From (R-E), it also immediately follows that

$$\text{for all } o \text{ occurring in } p,\ \mathcal{O}[\![o]\!]r = \mathcal{O}[\![o]\!]r' \tag{R-O}$$

Now we can prove the main statement by induction on the depth i of p.

Base Case ($i = 0$). p^0 has the form (e target : *expr* obl : o^*), thus the clause (S-4a) is used to determine $\mathcal{P}[\![p]\!]r$. The thesis then trivially follows from (R-E) and (R-O).

Inductive Case ($i = n + 1$). p^{n+1} is of the form $\{a$ target : *expr* policies : $(p^+)^n$ obl-p : o_p^* obl-d : $o_d^*\}$, thus the clause (S-4b) is used to determine $\mathcal{P}[\![p]\!]r$. By the induction hypothesis, for any p_j^k in $(p^+)^n$, with $k \leq n$, it holds that $\mathcal{P}[\![p_j^k]\!]r = \mathcal{P}[\![p_j^k]\!]r'$. This, due to the clauses (S-6a) and (S-6b), implies that $\mathcal{A}[\![a, (p^+)^n]\!]r = \mathcal{A}[\![a, (p^+)^n]\!]r'$, for any algorithm a. The thesis then follows from this fact and from (R-E) and (R-O). □

B.2 Proofs of results in Section 6

Theorem 6.1 (Total and Deterministic Constraint Semantics).

1. For all $c \in \mathit{Constr}$ and $r \in R$, there exists an $el \in (\mathit{Value} \cup 2^{\mathit{Value}} \cup \{\mathsf{error}, \bot\})$, such that $\mathcal{C}[\![c]\!]r = el$.

2. For all $c \in \mathit{Constr}$, $r \in R$ and $el, el' \in (\mathit{Value} \cup 2^{\mathit{Value}} \cup \{\mathsf{error}, \bot\})$, it holds that

$$\mathcal{C}[\![c]\!]r = el \ \wedge\ \mathcal{C}[\![c]\!]r = el' \ \Rightarrow\ el = el'\,.$$

Proof. Similarly to the proof of Theorem 5.1, the proof reduces to showing that $\mathcal{C}$ is a total and deterministic function. We proceed by structural induction on the syntax of c.

Base Case. If $c = v$, the thesis immediately follows since $\mathcal{C}[\![v]\!]r = v$; otherwise, i.e. $c = n$, we have $\mathcal{C}[\![n]\!]r = r(n)$ and the thesis follows because r is a total and deterministic function.

Inductive Case. It is not hard to believe that all the defining clauses of $\mathcal{C}$ are such that the conditions in the right hand side are mutually exclusive and cover all the necessary cases. For each different form that c can assume, the thesis then directly follows by the induction hypothesis. □

The proof of Theorem 6.2 relies on the following three auxiliary results.

Lemma B.1. *For all expr ∈ Expr and r ∈ R, it holds that*

$$\mathcal{E}[\![\mathit{expr}]\!]r = \mathcal{C}[\![\mathcal{T}_E\{\!|\mathit{expr}|\!\}]\!]r$$

Proof. We proceed by structural induction on the syntax of *expr* according to the translation rules of the clause (T-1).

(*expr* = n). Since $\mathcal{T}_E\{\!|n|\!\} = n$, the thesis follows because $\mathcal{E}[\![n]\!]r = r(n) = \mathcal{C}[\![n]\!]r$.

($expr = v$). Since $\mathcal{T}_E\{\!|v|\!\} = v$, the thesis follows because $\mathcal{E}[\![v]\!]r = v = \mathcal{C}[\![v]\!]r$.

($expr = \mathsf{not}(expr_1)$). Since $\mathcal{T}_E\{\!|expr|\!\} = \dot{\neg}\, \mathcal{T}_E\{\!|expr_1|\!\}$ and, by the induction hypothesis, $\mathcal{E}[\![expr_1]\!]r = \mathcal{C}[\![\mathcal{T}_E\{\!|expr_1|\!\}]\!]r$, the thesis follows due to the correspondence of the semantic clause of the operator $\dot{\neg}$ in Table 7 and that of the operator not in Table 5.

($expr = \mathsf{op}(expr_1, expr_2)$). Since $\mathcal{T}_E\{\!|expr|\!\} = \mathcal{T}_E\{\!|expr_1|\!\}\ \mathsf{getOp}(\mathsf{op})\ \mathcal{T}_E\{\!|expr_2|\!\}$ and, by the induction hypothesis, $\mathcal{E}[\![expr_1]\!]r = \mathcal{C}[\![\mathcal{T}_E\{\!|expr_1|\!\}]\!]r$ and $\mathcal{E}[\![expr_2]\!]r = \mathcal{C}[\![\mathcal{T}_E\{\!|expr_2|\!\}]\!]r$, the thesis follows due to the correspondence of the semantic clause of the expression operator op in Table 7 and that of the constraint operator $\mathsf{getOp}(\mathsf{op})$ in Table 5.

□

Lemma B.2. *For all $o \in Obligation$ and $r \in R$ it holds that*

$$\mathcal{O}[\![o]\!]r = io \quad \Leftrightarrow \quad \mathcal{C}[\![\mathcal{T}_{Ob}\{\!|o|\!\}]\!]r = \mathsf{true}$$

and

$$\mathcal{O}[\![o]\!]r = \mathsf{error} \quad \Leftrightarrow \quad \mathcal{C}[\![\mathcal{T}_{Ob}\{\!|o|\!\}]\!]r = \mathsf{false}$$

Proof. We only prove the ($\Rightarrow$) implication as the proof for the other direction proceeds in a specular way. Let $o = [t\ pepAct(expr^*)]$ with $expr^* = expr_1 \ldots expr_n$. By the clause (T-2), it is translated into the constraint

$$c \;=\; \textstyle\bigwedge_{expr_j \in expr^*} \neg\mathtt{isMiss}(\mathcal{T}_E\{\!|expr_j|\!\}) \wedge \neg\mathtt{isErr}(\mathcal{T}_E\{\!|expr_j|\!\})$$

We now proceed by case analysis on $\mathcal{O}[\![o]\!]r$.

($\mathcal{O}[\![o]\!]r = io$). We have to prove that $\mathcal{C}[\![c]\!]r = \mathsf{true}$. By the definition of $\mathcal{C}$, $\mathcal{C}[\![c]\!]r = \mathsf{true}$ corresponds to

$$\begin{array}{l}\forall j \in \{1, \ldots, n\}\ : \\ \quad \mathcal{C}[\![\neg\ \mathtt{isMiss}(\mathcal{T}_E\{\!|expr_j|\!\})]\!]r = \mathsf{true} \\ \quad\quad \wedge\ \mathcal{C}[\![\neg\ \mathtt{isErr}(\mathcal{T}_E\{\!|expr_j|\!\})]\!]r = \mathsf{true}\end{array}$$

According to the constraint semantics of $\neg$, `isMiss` and `isErr`, this corresponds to

$$\begin{array}{l}\forall j \in \{1, \ldots, n\}\ : \\ \quad \mathcal{C}[\![\mathcal{T}_E\{\!|expr_j|\!\}]\!]r \neq \bot \\ \quad\quad \wedge\ \ \mathcal{C}[\![\mathcal{T}_E\{\!|expr_j|\!\}]\!]r \neq \mathsf{error}\end{array}$$

By the hypothesis $\mathcal{O}[\![o]\!]r = io$ and the clauses (S-3a) and (S-2), we have

$$\mathcal{E}[\![expr^*]\!]r = \mathcal{E}[\![expr_1]\!]r \bullet \ldots \bullet \mathcal{E}[\![expr_n]\!]r = w_1\ \ldots w_n$$

where w_j stands for a literal value or a set of values. Thus, by Lemma B.1, we get that

$$\begin{array}{l}\forall j \in \{1, \ldots, n\}\ : \\ \quad \mathcal{C}[\![\mathcal{T}_E\{\!|expr_j|\!\}]\!]r = w_j \notin \{\bot, \mathsf{error}\}\end{array}$$

which proves the thesis.

($\mathcal{O}[\![o]\!]r = \mathsf{error}$). We have to prove that $\mathcal{C}[\![c]\!]r = \mathsf{false}$. By the definition of $\mathcal{C}$, $\mathcal{C}[\![c]\!]r = \mathsf{false}$ corresponds to

$$\begin{array}{l}\exists j \in \{1, \ldots, n\}\ : \\ \quad \mathcal{C}[\![\neg\ \mathtt{isMiss}(\mathcal{T}_E\{\!|expr_j|\!\})]\!]r = \mathsf{false} \\ \quad\quad \vee\ \mathcal{C}[\![\neg\ \mathtt{isErr}(\mathcal{T}_E\{\!|expr_j|\!\})]\!]r = \mathsf{false}\end{array}$$

According to the constraint semantics of $\neg$, `isMiss` and `isErr`, this corresponds to

$$\begin{aligned}&\exists j \in \{1,\ldots,n\} \ : \\ &\quad \mathcal{C}[\![\mathcal{T}_E\{\!|expr_j|\!\}]\!]r = \bot \ \vee \ \mathcal{C}[\![\mathcal{T}_E\{\!|expr_j|\!\}]\!]r = \mathsf{error}\end{aligned}$$

By the hypothesis $\mathcal{O}[\![o]\!]r = \mathsf{error}$ and the clauses (S-3a) and (S-2), we have

$$\begin{aligned}\mathcal{E}[\![expr^*]\!]r = \mathcal{E}[\![expr_1]\!]r \bullet \ldots \bullet \mathcal{E}[\![expr_n]\!]r \neq w^* \\ \Rightarrow \ \exists j \in \{1,\ldots,n\} : \mathcal{E}[\![expr_j]\!]r \in \{\bot, \mathsf{error}\}\end{aligned}$$

Thus, by Lemma B.1, we obtain that

$$\exists j \in \{1,\ldots,n\} \ : \ \mathcal{C}[\![\mathcal{T}_E\{\!|expr_j|\!\}]\!]r \in \{\bot, \mathsf{error}\}$$

which proves the thesis. □

Lemma B.3. *For all* $\mathsf{alg}_{\mathsf{all}} \in Alg$, $r \in R$ *and policies* $p_1,\ldots,p_s \in Policy$ *such that* $\forall\, i \in \{1,\ldots,s\}$: $\mathcal{P}[\![p_i]\!]r = \langle dec_i \ io_i^* \rangle \ \Leftrightarrow \ \mathcal{C}[\![\mathcal{T}_P\{\!|p_i|\!\} \downarrow_{dec_i}]\!]r = \mathsf{true}$, *it holds that*

$$\begin{aligned}&\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, p_1 \ \ldots \ p_s]\!]r = \langle dec \ io^* \rangle \ \Leftrightarrow \\ &\mathcal{C}[\![\mathcal{T}_A\{\!|\mathsf{alg}_{\mathsf{all}}, p_1 \ \ldots \ p_s|\!\} \downarrow_{dec}]\!]r = \mathsf{true}\end{aligned}$$

Proof. Since the considered algorithms use the all instantiation strategy, by the hypothesis and the clauses (S-6a) and (T-4), the thesis is equivalent to prove that

$$\begin{aligned}&\otimes\mathsf{alg}(\ldots \otimes \mathsf{alg}(\langle dec_1 \ io_1^* \rangle, \langle dec_2 \ io_2^* \rangle), \ldots, \langle dec_s \ io_s^* \rangle) \\ &\qquad\qquad = \langle dec \ io^* \rangle \\ &\qquad\qquad \Longleftrightarrow \\ &\mathcal{C}[\![\mathsf{alg}(\ldots \mathsf{alg}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\}), \ldots, \mathcal{T}_P\{\!|p_s|\!\}) \downarrow_{dec}]\!]r \\ &\qquad\qquad = \mathsf{true}\end{aligned}$$

The proof proceeds by case analysis on alg. In what follows, we only report the case of the p-over algorithm, as the other ones are similar and derive directly from the definitions in Appendix A

When $s = 1$, we have $\otimes\mathsf{p\text{-}over}(\mathcal{P}[\![p_1]\!]r) = \mathcal{P}[\![p_1]\!]r$ and $\mathsf{p\text{-}over}(\mathcal{T}_P\{\!|p_1|\!\}) = \mathcal{T}_P\{\!|p_1|\!\}$ by definition, hence the thesis directly follows from the hypothesis that $\mathcal{P}[\![p_1]\!]r = \langle dec_1 \ io_1^* \rangle \ \Leftrightarrow \ \mathcal{C}[\![\mathcal{T}_P\{\!|p_1|\!\} \downarrow_{dec_1}]\!]r = \mathsf{true}$. For the remaining cases, we proceed by induction on the number s of policies to combine.

Base Case ($s = 2$). We must prove that

$$\begin{aligned}&\otimes\mathsf{p\text{-}over}(\langle dec_1 \ io_1^* \rangle, \langle dec_2 \ io_2^* \rangle) = \langle dec \ io^* \rangle \\ &\qquad\qquad \Leftrightarrow \\ &\mathcal{C}[\![\mathsf{p\text{-}over}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\}) \downarrow_{dec}]\!]r = \mathsf{true}\,.\end{aligned}$$

For the sake of simplicity, in the following we omit the sequences of instantiated obligations, as their combination does not affect the decision *dec* returned by $\otimes$p-over. We proceed by case analysis on the decision *dec*.

($dec = \mathsf{permit}$). It follows that $dec_1 = \mathsf{permit}$ or $dec_2 = \mathsf{permit}$. Moreover, by definition we have $\mathsf{p\text{-}over}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\}) \downarrow_p = \mathcal{T}_P\{\!|p_1|\!\} \downarrow_p \ \vee \ \mathcal{T}_P\{\!|p_2|\!\} \downarrow_p$.

($dec = \mathsf{deny}$). It follows that $dec_1, dec_2 \in \{\mathsf{deny}, \mathsf{not\text{-}app}\}$. Moreover, by definition we have $\mathsf{p\text{-}over}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\}) \downarrow_d = (\mathcal{T}_P\{\!|p_1|\!\} \downarrow_d \ \wedge \mathcal{T}_P\{\!|p_2|\!\} \downarrow_d) \vee (\mathcal{T}_P\{\!|p_1|\!\} \downarrow_d \wedge \mathcal{T}_P\{\!|p_2|\!\} \downarrow_n) \vee (\mathcal{T}_P\{\!|p_1|\!\} \downarrow_n \wedge \mathcal{T}_P\{\!|p_2|\!\} \downarrow_d)$.

($dec = \mathsf{not\text{-}app}$). It follows that $dec_1 = dec_2 = \mathsf{not\text{-}app}$. Moreover, by definition we have $\mathsf{p\text{-}over}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\})\downarrow_n = \mathcal{T}_P\{\!|p_1|\!\}\downarrow_n \wedge \mathcal{T}_P\{\!|p_2|\!\}\downarrow_n$.

($dec = \mathsf{indet}$). It follows that $dec_1 = \mathsf{indet}$ or $dec_2 = \mathsf{indet}$ and $dec_1, dec_2 \neq \mathsf{permit}$. Moreover, by definition we have $\mathsf{p\text{-}over}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\})\downarrow_i = (\mathcal{T}_P\{\!|p_1|\!\}\downarrow_i \wedge \neg\mathcal{T}_P\{\!|p_2|\!\}\downarrow_p) \vee (\neg\mathcal{T}_P\{\!|p_1|\!\}\downarrow_p \wedge \mathcal{T}_P\{\!|p_2|\!\}\downarrow_i)$.

In any case, thesis follows from the hypothesis on $\mathcal{T}_P\{\!|p_i|\!\}$ and the definition of $\mathcal{C}$.

Inductive Case ($s = k + 1$). By the induction hypothesis the thesis holds for k policies, that is

$$\otimes\mathsf{alg}(\ldots \otimes \mathsf{alg}(\langle dec_1\ io_1^*\rangle, \langle dec_2\ io_2^*\rangle), \ldots, \langle dec_k\ io_k^*\rangle) = \langle dec\ io^*\rangle$$

$$\Longleftrightarrow$$

$$\mathcal{C}[\![\mathsf{alg}(\ldots \mathsf{alg}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\}), \ldots, \mathcal{T}_P\{\!|p_k|\!\})\downarrow_{dec}]\!]r = \mathsf{true}$$

The thesis then follows by repeating the case analysis on decision dec of the 'Base Case' once we replace $\langle dec_1\ io_1^*\rangle$, $\langle dec_2\ io_2^*\rangle$, $\mathcal{T}_P\{\!|p_1|\!\}$ and $\mathcal{T}_P\{\!|p_2|\!\}$ by $\langle dec'\ io'^*\rangle$, $\langle dec_s\ io_s^*\rangle$, $\mathsf{p\text{-}over}(\ldots \mathsf{p\text{-}over}(\mathcal{T}_P\{\!|p_1|\!\}, \mathcal{T}_P\{\!|p_2|\!\}), \ldots, \mathcal{T}_P\{\!|p_k|\!\})$ and $\mathcal{T}_P\{\!|p_s|\!\}$, respectively. □

THEOREM 6.2 [Policy Semantic Correspondence] For all $p \in \mathit{Policy}$ enclosing combining algorithms only using all as instantiation strategy, and $r \in R$, it holds that

$$\mathcal{P}[\![p]\!]r = \langle dec\ io^*\rangle \quad \Leftrightarrow \quad \mathcal{C}[\![\mathcal{T}_P\{\!|p|\!\}\downarrow_{dec}]\!]r = \mathsf{true}$$

Proof. The proof proceeds by induction on the depth i of p.

Base Case ($i = 0$). This means that p is of the form $(e\ \ \mathsf{target} : expr\ \ \mathsf{obl} : o^*)$. We proceed by case analysis on dec.

($dec = \mathsf{permit}$). By the clause (S-4a), it follows that

$$\mathcal{E}[\![expr]\!]r = \mathsf{true} \ \wedge\ \mathcal{O}[\![o^*]\!]r = io^*$$

Thus, by Lemma B.1, it follows that

$$\mathcal{C}[\![\mathcal{T}_E\{\!|expr|\!\}]\!]r = \mathsf{true}$$

and, by Lemma B.2 and the clause (T-2), it follows that

$$\mathcal{C}[\![\mathcal{T}_{Ob}\{\!|o^*|\!\}]\!]r = \mathsf{true}$$

On the other hand, by the clause (T-3a), we have that

$$\mathcal{T}_P\{\!|(\mathsf{permit}\ \ \mathsf{target} : expr\ \ \mathsf{obl} : o^*)|\!\}\downarrow_p = \mathcal{T}_E\{\!|expr|\!\} \wedge \mathcal{T}_{Ob}\{\!|o^*|\!\}$$

Hence, by the definition of $\mathcal{C}$, we can conclude that

$$\mathcal{C}[\![\mathcal{T}_P\{\!|(\mathsf{permit}\ \ \mathsf{target} : expr\ \ \mathsf{obl} : o^*)|\!\}\downarrow_p]\!]r = \mathcal{C}[\![\mathcal{T}_E\{\!|expr|\!\}]\!]r \wedge \mathcal{C}[\![\mathcal{T}_{Ob}\{\!|o^*|\!\}]\!]r = \mathsf{true} \wedge \mathsf{true} = \mathsf{true}$$

which proves the thesis.

($dec = \mathsf{deny}$). We omit the proof since it proceeds like the previous case.

($dec = \mathsf{not\text{-}app}$). By the clause (S-4a), it follows that

$$\mathcal{E}[\![expr]\!]r = \mathsf{false} \ \vee \ \mathcal{E}[\![expr]\!]r = \bot$$

By the clause (T-3a), we have that

$$\mathcal{T}_P\{\!|(e \ \ \mathsf{target} : expr \ \ \mathsf{obl} : o^* \,)|\!\}\downarrow_n = \neg\, \mathcal{T}_E\{\!|expr|\!\}$$

Hence, the thesis directly follows by Lemma B.1 and the definition of $\mathcal{C}$.

($dec = \mathsf{indet}$). By the clause (S-4a), the `otherwise` condition holds, that is

$$\begin{array}{l}\neg(\mathcal{E}[\![expr]\!]r = \mathsf{true} \ \wedge \ \mathcal{O}[\![o^*]\!]r = io^*) \\ \wedge\, \neg(\mathcal{E}[\![expr]\!]r = \mathsf{false} \ \vee \ \mathcal{E}[\![expr]\!]r = \bot)\end{array}$$

By applying standard boolean laws and reasoning on function codomains, this condition can be rewritten as follows

$$\begin{array}{l}\neg(\mathcal{E}[\![expr]\!]r = \mathsf{true} \ \wedge \ \mathcal{O}[\![o^*]\!]r = io^*) \\ \quad \wedge\, \neg(\mathcal{E}[\![expr]\!]r = \mathsf{false} \ \vee \ \mathcal{E}[\![expr]\!]r = \bot) \\ = (\mathcal{E}[\![expr]\!]r \neq \mathsf{true} \vee \mathcal{O}[\![o^*]\!]r = \mathsf{error}) \\ \quad \wedge\, (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{false}, \bot\}) \\ = \mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{false}, \bot\} \\ \quad \wedge\, \mathcal{O}[\![o^*]\!]r = \mathsf{error}) \\ = \mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \vee \\ \quad (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \wedge \mathcal{O}[\![o^*]\!]r = \mathsf{error}) \vee \\ \quad (\mathcal{E}[\![expr]\!]r = \mathsf{true} \wedge \mathcal{O}[\![o^*]\!]r = \mathsf{error}) \\ = \mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\ \quad \vee (\mathcal{E}[\![expr]\!]r = \mathsf{true} \wedge \mathcal{O}[\![o^*]\!]r = \mathsf{error})\end{array}$$

On the other hand, by the clause (T-3a), we have that

$$\begin{array}{l}\mathcal{T}_P\{\!|(e \ \ \mathsf{target} : expr \ \ \mathsf{obl} : o^* \,)|\!\}\downarrow_i = \\ \quad \neg\ (\mathtt{isBool}(\mathcal{T}_E\{\!|expr|\!\}) \vee \mathtt{isMiss}(\mathcal{T}_E\{\!|expr|\!\})) \\ \quad \vee\ (\mathcal{T}_E\{\!|expr|\!\} \wedge \neg\, \mathcal{T}_{Ob}\{\!|o^*|\!\})\end{array}$$

The thesis then follows by Lemmas B.1 and B.2 and the definition of $\mathcal{C}$.

Inductive Case ($i = k + 1$). p is of the form $\{\mathsf{alg}_{\mathsf{all}} \ \ \mathsf{target} : expr \ \ \mathsf{policies} : (p^+)^k \, \mathsf{obl\text{-}p} : o_p^* \ \ \mathsf{obl\text{-}d} : o_d^* \,\}$. We proceed by case analysis on dec.

($dec = \mathsf{permit}$). By the clause (S-4b), it follows that

$$\begin{array}{l}\mathcal{E}[\![expr]\!]r = \mathsf{true} \\ \wedge\, \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle \mathsf{permit} \ \ io_1^* \rangle \\ \wedge\, \mathcal{O}[\![o_p^*]\!]r = io_2^*\end{array}$$

Thus, by Lemma B.1, it follows that

$$\mathcal{E}[\![expr]\!]r = \ \mathcal{C}[\![\mathcal{T}_E\{\!|expr|\!\}]\!]r = \mathsf{true}$$

and, by Lemma B.2 and the clause (T-2), it follows that

$$\mathcal{C}[\![\mathcal{T}_{Ob}\{\!|o_p^*|\!\}]\!]r = \mathsf{true}$$

Since by the induction hypothesis, for all p_i^h in $(p^+)^k$ with $h \leq k$, it holds that

$$\mathcal{P}[\![p_i^h]\!]r = \langle dec_i \ io^* \rangle \quad \Leftrightarrow \quad \mathcal{C}[\![\mathcal{T}_P\{\!|p_i^h|\!\}\downarrow_{dec_i}]\!]r = \mathsf{true}$$

then, by Lemma B.3, it follows that

$$\mathcal{T}_A\{\!|\mathsf{alg}_{\mathsf{all}}, (p^+)^k|\!\}\downarrow_p = \mathsf{true}$$

On the other hand, by the clause (T-3b), we have that

$$\begin{aligned}
&\mathcal{T}_P\{\!|\{\mathsf{alg}_{\mathsf{all}}\ \mathsf{target} : expr\ \mathsf{policies} : (p^+)^k \\
&\qquad\qquad\qquad\qquad \mathsf{obl\text{-}p} : o_p^*\ \ \mathsf{obl\text{-}d} : o_d^*\ \}|\!\}\downarrow_p \\
&= \mathcal{T}_E\{\!|expr|\!\} \wedge\ \mathcal{T}_A\{\!|\mathsf{alg}_{\mathsf{all}}, (p^+)^k|\!\}\downarrow_p \wedge\ \mathcal{T}_{Ob}\{\!|o_p^*|\!\}
\end{aligned}$$

Hence, by the definition of $\mathcal{C}$, we can conclude that

$$\begin{aligned}
&\mathcal{C}[\![\mathcal{T}_P\{\!|\{\mathsf{alg}_{\mathsf{all}}\ \mathsf{target} : expr\ \mathsf{policies} : (p^+)^k \\
&\qquad\qquad\qquad\qquad \mathsf{obl\text{-}p} : o_p^*\ \ \mathsf{obl\text{-}d} : o_d^*\ \}|\!\}\downarrow_p]\!]r \\
&= \mathcal{C}[\![\mathcal{T}_E\{\!|expr|\!\}]\!]r \wedge \mathcal{C}[\![\mathcal{T}_A\{\!|\mathsf{alg}_{\mathsf{all}}, (p^+)^k|\!\}\downarrow_p]\!]r \\
&\quad \wedge\ \mathcal{C}[\![\mathcal{T}_{Ob}\{\!|o_p^*|\!\}]\!]r \\
&= \mathsf{true} \wedge \mathsf{true} \wedge \mathsf{true} = \mathsf{true}
\end{aligned}$$

which proves the thesis.

($dec = \mathsf{deny}$). We omit the proof since it proceeds like the previous case.

($dec = \mathsf{not\text{-}app}$). By the clause (S-4b), it follows that

$$\begin{aligned}
&\mathcal{E}[\![expr]\!]r = \mathsf{false} \vee\ \mathcal{E}[\![expr]\!]r = \bot \\
&\vee\ (\mathcal{E}[\![expr]\!]r = \mathsf{true}\ \wedge\ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \mathsf{not\text{-}app})
\end{aligned}$$

By the clause (T-3b), we have that

$$\begin{aligned}
&\mathcal{T}_P\{\!|\{\mathsf{alg}_{\mathsf{all}}\ \mathsf{target} : expr\ \mathsf{policies} : (p^+)^k \\
&\qquad\qquad\qquad\qquad \mathsf{obl\text{-}p} : o_p^*\ \ \mathsf{obl\text{-}d} : o_d^*\ \}|\!\}\downarrow_n \\
&= \neg\ \mathcal{T}_E\{\!|expr|\!\} \vee (\mathcal{T}_E\{\!|expr|\!\} \wedge \mathcal{T}_A\{\!|\mathsf{alg}_{\mathsf{all}}, (p^+)^k|\!\}\downarrow_n)
\end{aligned}$$

The thesis then directly follows by Lemmas B.1 and B.3, due to the induction hypothesis and the definition of $\mathcal{C}$.

($dec = \mathsf{indet}$). By the clause (S-4b), the `otherwise` condition holds, that is

$$\begin{aligned}
&\neg(\mathcal{E}[\![expr]\!]r = \mathsf{true}\ \wedge\ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e\ \ io_1^* \rangle \\
&\quad \wedge\ (\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e\ \ io_1^* \rangle \Longrightarrow \mathcal{O}[\![o_e^*]\!]r = io_2^*) \\
&\wedge\ \neg(\mathcal{E}[\![expr]\!]r = \mathsf{false} \vee\ \mathcal{E}[\![expr]\!]r = \bot \\
&\quad \vee\ (\mathcal{E}[\![expr]\!]r = \mathsf{true}\ \wedge\ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \mathsf{not\text{-}app}))
\end{aligned} \tag{C}$$

where, to recall the connection between the effect returned by the combining algorithm and the sequence of obligations that is instantiated, we exploit the tautology

$$\begin{aligned}
&\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e\ \ io_1^* \rangle \wedge \mathcal{O}[\![o_e^*]\!]r = io_2^* \\
&= \\
&\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e\ \ io_1^* \rangle \\
&\wedge(\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e\ \ io_1^* \rangle \Longrightarrow \mathcal{O}[\![o_e^*]\!]r = io_2^*)
\end{aligned}$$

By applying standard boolean laws and reasoning on function codomains, the Condition (C) above can be rewritten as follows

$$
\begin{array}{l}
\neg(\mathcal{E}[\![expr]\!]r = \mathsf{true} \ \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \\
\quad \wedge \ (\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \Longrightarrow \mathcal{O}[\![o_e^*]\!]r = io_2^*) \\
\wedge \ \neg(\mathcal{E}[\![expr]\!]r = \mathsf{false} \vee \ \mathcal{E}[\![expr]\!]r = \bot \\
\quad \vee \ (\mathcal{E}[\![expr]\!]r = \mathsf{true} \ \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \mathsf{not\text{-}app})) \\
= \\
(\mathcal{E}[\![expr]\!]r \neq \mathsf{true} \ \vee \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \in \{\mathsf{not\text{-}app}, \mathsf{indet}\} \\
\quad \vee \ (\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \wedge \mathcal{O}[\![o_e^*]\!]r = \mathsf{error})) \\
\wedge \ (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{false}, \bot\} \\
\quad \wedge (\mathcal{E}[\![expr]\!]r \neq \mathsf{true} \ \vee \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \neq \mathsf{not\text{-}app})) \\
= \\
(\mathcal{E}[\![expr]\!]r \neq \mathsf{true} \ \vee \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \in \{\mathsf{not\text{-}app}, \mathsf{indet}\} \\
\quad \vee \ (\mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \wedge \mathcal{O}[\![o_e^*]\!]r = \mathsf{error})) \\
\wedge (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\quad \vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{false}, \bot\} \wedge \\
\quad \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \neq \mathsf{not\text{-}app})) \\
= \\
\quad \mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\quad \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \neq \mathsf{not\text{-}app}) \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\quad \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \in \{\mathsf{not\text{-}app}, \mathsf{indet}\}) \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{false}, \bot\} \\
\quad \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \neq \mathsf{not\text{-}app} \\
\quad \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \in \{\mathsf{not\text{-}app}, \mathsf{indet}\}) \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \\
\quad \wedge \ \mathcal{O}[\![o_e^*]\!]r = \mathsf{error}) \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{false}, \bot\} \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r \neq \mathsf{not\text{-}app} \\
\quad \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \wedge \mathcal{O}[\![o_e^*]\!]r = \mathsf{error}) \\
= \\
\quad \mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{false}, \bot\} \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \mathsf{indet}) \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{false}, \bot\} \\
\quad \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \wedge \mathcal{O}[\![o_e^*]\!]r = \mathsf{error}) \\
= \\
\quad \mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \mathsf{indet}) \\
\vee (\mathcal{E}[\![expr]\!]r = \mathsf{true} \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \mathsf{indet}) \\
\vee (\mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\quad \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \wedge \mathcal{O}[\![o_e^*]\!]r = \mathsf{error}) \\
\vee (\mathcal{E}[\![expr]\!]r = \mathsf{true} \\
\quad \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \wedge \mathcal{O}[\![o_e^*]\!]r = \mathsf{error}) \\
= \\
\end{array}
$$

$$
\begin{array}{l}
\quad \mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\vee(\mathcal{E}[\![expr]\!]r = \mathsf{true} \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \mathsf{indet}) \\
\vee(\mathcal{E}[\![expr]\!]r = \mathsf{true} \wedge \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle e \quad io_1^* \rangle \\
\quad \wedge \mathcal{O}[\![o_e^*]\!]r = \mathsf{error}) \\
= \\
\quad \mathcal{E}[\![expr]\!]r \notin \{\mathsf{true}, \mathsf{false}, \bot\} \\
\vee(\mathcal{E}[\![expr]\!]r = \mathsf{true} \ \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \mathsf{indet} \ ) \\
\vee(\mathcal{E}[\![expr]\!]r = \mathsf{true} \ \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle \mathsf{permit} \quad io_1^* \rangle \\
\quad \wedge \mathcal{O}[\![o_p^*]\!]r = \mathsf{error}) \\
\vee(\mathcal{E}[\![expr]\!]r = \mathsf{true} \ \wedge \ \mathcal{A}[\![\mathsf{alg}_{\mathsf{all}}, (p^+)^k]\!]r = \langle \mathsf{deny} \quad io_1^* \rangle \\
\quad \wedge \mathcal{O}[\![o_d^*]\!]r = \mathsf{error})
\end{array}
$$

where the last step exploits the fact that $e \in \{\mathsf{permit}, \mathsf{deny}\}$.
On the other hand, by the clause (T-3b), we have that

$$
\begin{array}{l}
\mathcal{T}_P\{\!|\{\mathsf{alg}_{\mathsf{all}} \ \mathsf{target} : expr \ \mathsf{policies} : (p^+)^k \\
\qquad\qquad\qquad\qquad\qquad \mathsf{obl\text{-}p} : o_p^* \quad \mathsf{obl\text{-}d} : o_d^* \ \}|\!\} \downarrow_i = \\
\neg \ (\mathtt{isBool}(\mathcal{T}_E\{\!|expr|\!\}) \ \vee \ \mathtt{isMiss}(\mathcal{T}_E\{\!|expr|\!\})) \\
\vee \ (\mathcal{T}_E\{\!|expr|\!\} \wedge \ \mathcal{T}_A\{\!|a, (p^+)^k|\!\} \downarrow_i) \\
\vee \ (\mathcal{T}_E\{\!|expr|\!\} \wedge \ \mathcal{T}_A\{\!|a, (p^+)^k|\!\} \downarrow_p \wedge \neg \ \mathcal{T}_{Ob}\{\!|o_p^*|\!\} \ ) \\
\vee \ (\mathcal{T}_E\{\!|expr|\!\} \wedge \ \mathcal{T}_A\{\!|a, (p^+)^k|\!\} \downarrow_d \wedge \neg \ \mathcal{T}_{Ob}\{\!|o_d^*|\!\} \ )
\end{array}
$$

The thesis then follows by Lemmas B.1, B.2 and B.3, due to the induction hypothesis and the definition of $\mathcal{C}$.

□